International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
Peer Reviewed, Referred & Indexed Journal
ISSN: 3068-272X www.ijdim.com Original Research Paper

Privacy-Preserving Generative AI in Healthcare Systems Using Federated Learning Approaches

Rajesh Poojari
Independent Researcher, USA.

Abstract- The research paper focuses on the mechanism of introducing Federated Learning alongside privacy-saving strategies in Generative Artificial Intelligence to healthcare applications. The study analyses the privacy protection/model accuracy trade-off by implementing Differential Privacy and Secure Aggregation. The synthetic datasets were applied to the model to train in five rounds and included many federated clients, that improved its accuracy by 49% to 60%. The findings suggest that Federated Learning has the potential to improve the performance of AI and preserve the privacy of data at the same time. Other challenges covered in the study include mode collapse and privacy-utility trade-offs, and recommended solutions to achieve the efficient and secure healthcare AI models.

Keywords: Federated Learning, privacy-saving strategies, Generative Artificial Intelligence, Healthcare applications, privacy protection, model accuracy trade-off, Differential Privacy, Secure Aggregation, synthetic datasets,

Received: 10-01-2026 Accepted: 18-02-2026 Published: 25-02-2026

I. INTRODUCTION

The development of a privacy-sensitive generative AI model in the medical domain based on federated learning (FL) makes sure that the model training is performed securely and without sending sensitive information to multiple institutions. This approach blends privacy promoting approaches like differential privacy, homomorphic encryption and secure aggregation, thereby weighing the need to have excellent predictive accuracy with an overarching strong privacy protection [1]. The use of mechanisms of this type is critical to developing AI-based healthcare solutions and ensuring patient confidentiality.

Problem statement

The problem to be considered in this study is the problem of ensuring data confidentiality when implementing generative AI in clinical sites [2]. The federated learning approaches enable the development of models that work in collaboration with a high number of healthcare facilities without passing any sensitive health data, thereby balancing out the two-fold needs of privacy and the effectiveness of AI-based healthcare solutions.

Research Aim

This research aims to explore the combination of federated learning and privacy-conserving methods in generative AI and expects to increase both data and model safety and healthcare system performance.

Research Objective

The particular objectives of the research are to:

- ***To review how federated learning can be applied to healthcare infrastructures;***
- ***To consider privacy-saving techniques including: differential privacy and secure aggregation;***
- ***To compare trade-offs of the privacy protection with the model accuracy;***
- ***To suggest a federated generative AI with increased privacy and specialized in healthcare.***

Novel Contribution

The proposed work advances a new concept of combining federated learning with enhanced privacy-enhancing techniques (like differential privacy and secure aggregation) to create safe and efficient generative artificial intelligence models in healthcare. The framework has concern over privacy issues of major importance and it retains the model performance and scalability.

II. LITERATURE REVIEW

1. Federation Learning in Healthcare Introduction

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
Peer Reviewed, Referred & Indexed Journal
ISSN: 3068-272X www.ijdim.com Original Research Paper

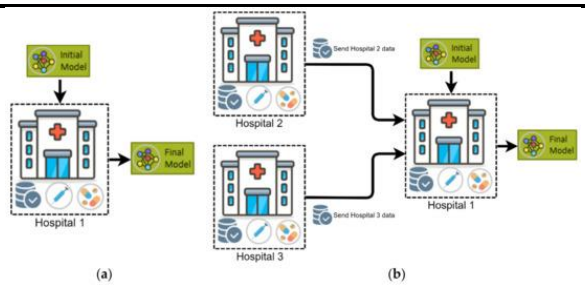

Fig. 1. Federated Learning (FL) In Healthcare

Federated Learning (FL) is a decentralized machine learning model in which two or more devices or institutions jointly train a model by avoiding the sensitive data distributors [3]. The method is of particular significance to healthcare, where the data can be stored at different institutions, including hospitals, clinics, and research facilities. Federated Learning provides the ability to comply with privacy laws like HIPAA and GDPR as the data about the patients remains in the local area [4]. The model has proved essential in healthcare contexts when used in diagnosis, predictive analytics, and decision support systems because, in most cases, privacy is a limiting factor to using centralized machine learning models. Federated Learning further tackles the problem of privacy because it provides an opportunity to keep the data on the local locations, and only updates on the models are sent [5]. This is essential in medical care, where the most important thing is data security, and the sharing of sensitive information of the patients can result in breaches. Federated Learning provides the potential to develop collaborative models without exposing data privacy that can be seen as a great solution to the problems of AI-based applications in healthcare.

2. Healthcare AI and Privacy-Preserving Techniques

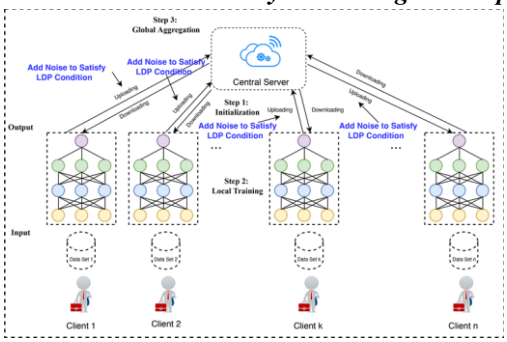

Fig. 2. Workflow of federated learning with local differential privacy

Privacy-preserving methods are the key to the healthcare privacy of patient information and the use of machine learning applications. The commonly used methods include Differential Privacy and Secure Aggregation [6]. ***Differential Privacy (DP)*** is a method that helps to preserve the individual data through the introduction of noise to the data update or model updates, thus causing the data to be indistinguishable as part of the aggregate result [7]. It also makes sure that the outputs of the model do not disclose information regarding the data of a particular participant. Differential Privacy has the potential to keep sensitive patient data safe in healthcare, even during training on data across several hospitals or clinics.

Another method that improves the privacy of Federated Learning is ***Secure Aggregation***. It enables the secure aggregation of multiple participant model updates without exposing any data of a participant [8]. This is to ensure that a server that gathers the updates cannot have access to the individual data points and hence, privacy is preserved in the course of the aggregation process.

3. Applicability of Federated Learning to Healthcare Systems

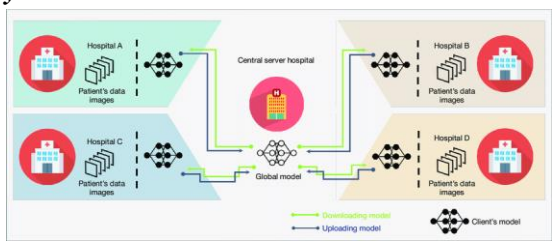

Fig. 3. Application of Federated Learning (FL) In Healthcare

Federated Learning has been of significant interest to the healthcare sector as it is capable of managing the privacy issue and still use the data to train the model. In the domain of healthcare, there is a tendency to spread the data across multiple establishments, and it is hard to centralize the data and use it to train the models based on large datasets [9]. Federated Learning enables such institutions to train models together without sending patient-sensitive information to a central server.

Federated Learning has also been used in medical imaging where it was used to train AI models to perform tasks like lung cancer detection on CT scans

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
Peer Reviewed, Referred & Indexed Journal
ISSN: 3068-272X www.ijdim.com Original Research Paper

[10]. This type of decentralized method will allow healthcare organizations to participate in model development without exchanging patient data. This method of collaboration enhances the generalizability of models and makes sure that the AI system is exposed to a wide range of information, which is crucial to the sound performance of models. Federated Learning has also been used in other medical projects like ***electronic health records (EHRs***) and genomics [11]. Through Federated Learning, health care organizations can have the ability to collaborate in training models aimed at predicting patient outcomes, risk detection, and treatment prescriptions, without jeopardizing patient privacy.

4. Federated Learning Models Privacy-Accuracy Trade-Offs

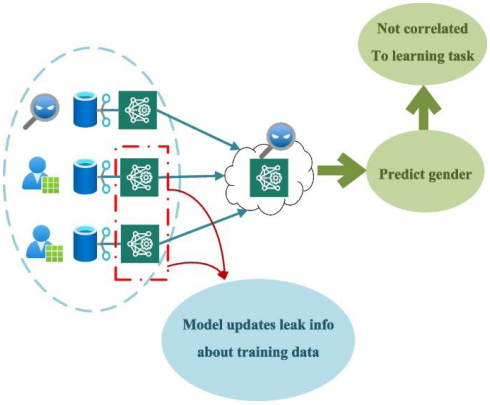

Fig. 4. Implementing security and privacy in federated learning

The main issue in Federated Learning is to strike a balance between the privacy of the information and the precision of the models. Although privacy-sensitive methods like Differential privacy and Secure Aggregation enhance the level of privacy in data, they may add noise in the model updates, thereby reducing the model performance [12]. Privacy-accuracy trade-off is the trade-off between privacy and model accuracy. Using more privacy such as one can obtain with Differential Privacy, the model could lose its precision as it loses information. This balance is paramount in the field of healthcare where accuracy of the models adopted is crucial in making correct predictions and recommendations to treatment [13].

In order to resolve this problem, different research works have discussed methods of maximizing the privacy-accuracy trade-off. The studies indicate that by properly setting privacy controls, one can have a compromise that is satisfactory so that privacy is not too much compromised but at the same time, the model does not perform poorly [14]. Using an appropriate amount of noise being added, healthcare organizations may ensure that their models are accurate and comply with privacy laws at the same time.

5. A Privacy-Enhanced Federated Generated AI Framework in Healthcare

A privacy-enhanced Federated Learning AI model may be suggested to healthcare systems to reduce the issues of privacy-accuracy trade-offs. The framework would integrate the merits of Federated Learning, privacy-offering methods and strategies, including Differential Privacy and Secure Aggregation, and Generative AI models such as Generative Adversarial Networks (GANs) [15]. It is possible to generate synthetic healthcare data using generative AI models that do not require sensitive data to simulate the statistical characteristics of real healthcare data. The Federated Learning models can be augmented with these synthetic datasets to enhance the performance of models without jeopardizing privacy [16]. This would allow the healthcare institutions to cumulatively train AI models using a larger and more diverse range of data, which would more effectively improve model accuracy and generalizability and protect patient privacy. Federated Learning may be used to acknowledge the privacy issues of the healthcare system, while still relying on the capabilities of machine learning through the combination of both approaches. This framework would enable institutions to develop AI models on healthcare applications in collaboration, without the institutions sharing sensitive patient information [17]. Privacy-preserving solutions, including Differential Privacy, are also applied to this framework, therefore, safeguarding the privacy of individual patient data, within the process.

Literature Gap

The privacy-preserving methods and decentralized model training used are mostly because the literature on Federated Learning in healthcare largely centers on the aforementioned methods. Nonetheless, the gap between the application of Generative AI and Federated Learning models to improve privacy is substantial. Also, further studies should be conducted to optimize the privacy-accuracy trade-off of

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
Peer Reviewed, Referred & Indexed Journal
ISSN: 3068-272X www.ijdim.com Original Research Paper

applications in the healthcare context and assess the performance of synthetic data generation in terms of model accuracy. Plugging these holes would further enhance the security of models, their generalization, and their feasible implementation in health care systems.

III. METHODOLOGY

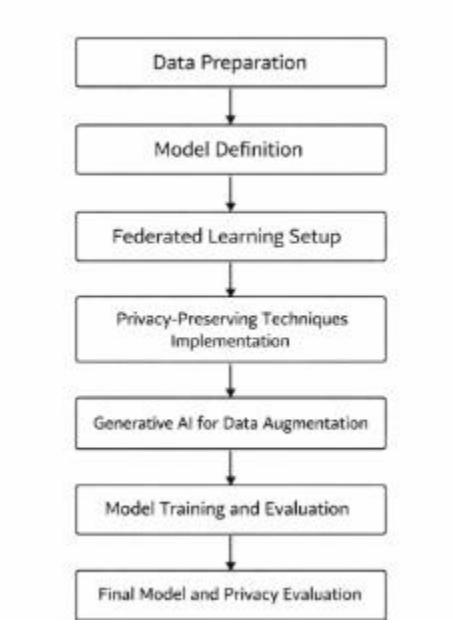

Fig 5. Research Flow Diagram

The approach to deploying a Federated Learning architecture and privacy-aware-techniques in healthcare consists of several important steps, such as preparing data, training the model, adding privacy-attractive features, testing and implementing privacy-friendly Generative AI models [18]. This will make use of Python software like TensorFlow, PyTorch, and Scikit-learn to make it easier to train models, and privacy-preserving algorithms including Differential Privacy and Secure Aggregation.

A. Data Preparation

The pre-treatment of data is an important initial stage of Federated learning. This information has to be distributed to several institutions or devices, which must also be local to the healthcare provider. The preprocessing procedure can contain the missing values, data normalization, and coding of categorical variables. Data preprocessing is an important process to achieve data quality and compatibility to machine learning process [19]. Treatment of missing values can be implemented using pandas filling and ordinal functions such as fillna, which is used to fill in the missing values, and ***dropna(),*** which is used to remove the incomplete rows. Normalization of data makes features similar to one another, and this becomes important when dealing with an algorithm that is sensitive to features [20]. It can be achieved by normalizing the features with ***sklearn.preprocessing.StandardScaler() or MinMaxScaler(),*** which are used respectively to normalize and standardize the features. Categorical variables will need encoding to convert categorical data to numeric form and this encoding may be accomplished through the encoding of categorical data by such methods as ***pandas.get_dummies() or sklearn.preprocessing.LabelEncoder().***

B. Model Definition

A machine learning model likeTensorFlow or PyTorch is used to define the model. The model of healthcare predictions may be a simple neural network:

$$\hat{y}i = f(Xi; \theta) \text{ ----- (1)}$$

Where:

- $\hat{y}i$ is the predicted output for the data
- Xi is the input data from institution i
- θ is the set of model parameters

C. Federated Learning Algorithm

In Federated Learning, the model is trained on many local datasets, and the updates related to models are sent to the central server instead of the actual data [21]. The model training algorithm is based on the iterative scheme of the Stochastic Gradient Descent (SGD):

$$\theta^{t+1} = \theta^t - \eta \nabla L(\theta^t, Xi, yi) \text{ ----- (2)}$$

Where:

- θ^t is the model's parameter at iteration t
- η is the learning rate
- ∇L is the gradient of the loss function L
- Xi and yi are the local data inputs and labels for institution i

D. Privacy-Preserving Techniques

Differential Privacy (DP): Differential Privacy in Federated Learning guarantees that the presence of an individual data point does not have a large impact on the output of the model [22]. Gradients are determined as noisy when updating the model to ensure privacy:

$$\widehat{\nabla Li} = \nabla Li + N(0, \sigma^2) \text{ ----- (3)}$$

Where:

- $\widehat{\nabla Li}$ is the differentially private gradient

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
Peer Reviewed, Referred & Indexed Journal
ISSN: 3068-272X www.ijdim.com Original Research Paper

- $N(0, \sigma^2)$ is Gaussian noise with mean 0 and variance σ2
- ∇Li is the local gradient

Secure Aggregation: Secure aggregation guarantees that the server not be able to find out the contribution of particular data [23]. Secure Aggregation can be mathematically represented as:

$$\hat{\theta}agg = \sum_{i=1}^{N} \quad \theta i \text{ ----- (4)}$$

Where:

- $\hat{\theta}agg$ is the aggregated model parameter
- θi is the model parameter from institution i
- N is the number of institutions or participants in the Federated Learning system

E. Training the Generative AI Model

In order to maximize privacy, a ***Generative Adversarial Network (GAN)*** may be applied to produce synthetic healthcare data. The GAN comprises a generator GGG and a discriminator DDD, which are trained in a minimax game:

$$minG\,.\,maxD\,.\,E_{\quad x\sim pdata}[logD(x)] + E_{\quad z\sim pz}$$
$$[log(1 - D(G(z)))] \text{ ----- (5)}$$

Where:

- G(z) is the generator mapping random noise z to data samples
- D(x) is the discriminator's estimate of whether x is a real data point or generated
- Pdata is the real data distribution
- Pz is the distribution of random noise

This works so as to create artificial healthcare data to supplement the Federated Learning process, which increases model generalizability without the loss of privacy.

F. Model Evaluation

The model is tested on a model on a privacy-preserving set of measures, such as model accuracy and privacy leakage and performance trade-offs. Also, Privacy Leakage is quantifiable by determining the vulnerability of the membership inference attack that tests if an attacker can determine whether a given data point is included in the training set.

IV. RESULT AND DISCUSSION

Model Performance in Federated Learning

```
import tensorflow as tf
import numpy as np
import collections
def create_keras_model():
    model = tf.keras.Sequential([
        tf.keras.layers.Dense(64, activation='relu'),
        tf.keras.layers.Dense(1, activation='sigmoid')
    ])
    return model
def create_mock_data():
    client_data = []
    for _ in range(5):
        x = np.random.rand(100, 1)
        y = (x > 0.5).astype(np.float32)
        client_data.append(collections.OrderedDict(
            x=tf.convert_to_tensor(x, dtype=tf.float32),
            y=tf.convert_to_tensor(y, dtype=tf.float32)
        ))
    return client_data
def client_train(model, client_data, epochs=1, batch_size=32):
    model.compile(optimizer='adam', loss='binary_crossentropy',
metrics=['accuracy'])
    model.fit(client_data['x'], client_data['y'], epochs=epochs,
batch_size=batch_size)
    return model.get_weights()
def federated_learning_round(global_model, federated_data, rounds=5):
    for round_num in range(rounds):
        client_weights = []
        for client_data in federated_data:
            local_model = create_keras_model()
            local_weights = client_train(local_model, client_data)
            client_weights.append(local_weights)
        if client_weights:
            num_params = len(client_weights[0])
            averaged_global_weights = []
            for param_idx in range(num_params):
                param_tensors_for_averaging = [client_w[param_idx] for
client_w in client_weights]

                averaged_param = np.mean(param_tensors_for_averaging,
axis=0)
                averaged_global_weights.append(averaged_param)
            global_model.set_weights(averaged_global_weights)
        else:
            print("No client weights collected for averaging.")
            continue
        loss, accuracy = global_model.evaluate(federated_data[0]['x'],
federated_data[0]['y'])
        print(f"Round {round_num + 1}: Loss = {loss:.4f}, Accuracy =
{accuracy:.4f}")
    return global_model
global_model = create_keras_model()
federated_data = create_mock_data()
global_model.build(input_shape=(None, federated_data[0]['x'].shape[1]))
global_model.compile(optimizer='adam', loss='binary_crossentropy',
metrics=['accuracy'])
global_model = federated_learning_round(global_model, federated_data,
rounds=5)
final_loss, final_accuracy =
global_model.evaluate(federated_data[0]['x'], federated_data[0]['y'])
print(f"Final model evaluation: Loss = {final_loss:.4f}, Accuracy =
{final_accuracy:.4f}")
```

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
Peer Reviewed, Referred & Indexed Journal
ISSN: 3068-272X www.ijdim.com Original Research Paper

```
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 11ms/step - accuracy: 0.5283 - loss: 0.7177
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 10ms/step - accuracy: 0.5473 - loss: 0.7121
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 10ms/step - accuracy: 0.6663 - loss: 0.6919
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 11ms/step - accuracy: 0.3902 - loss: 0.6947
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 10ms/step - accuracy: 0.5297 - loss: 0.7241
4/4 ━━━━━━━━━━━━━━━━━━━━ 0s 10ms/step - accuracy: 0.5252 - loss: 0.6974
Round 1: Loss = 0.6980, Accuracy = 0.4900
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 10ms/step - accuracy: 0.4112 - loss: 0.7073
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 18ms/step - accuracy: 0.4923 - loss: 0.6551
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 13ms/step - accuracy: 0.3585 - loss: 0.6973
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 10ms/step - accuracy: 0.4819 - loss: 0.7072
4/4 ━━━━━━━━━━━━━━━━━━━━ 2s 10ms/step - accuracy: 0.4915 - loss: 0.7148
4/4 ━━━━━━━━━━━━━━━━━━━━ 0s 10ms/step - accuracy: 0.4538 - loss: 0.6985
Round 2: Loss = 0.6991, Accuracy = 0.4000
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 10ms/step - accuracy: 0.4571 - loss: 0.7255
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 10ms/step - accuracy: 0.5483 - loss: 0.7002
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 10ms/step - accuracy: 0.4684 - loss: 0.7128
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 12ms/step - accuracy: 0.4383 - loss: 0.7416
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 10ms/step - accuracy: 0.5161 - loss: 0.6808
4/4 ━━━━━━━━━━━━━━━━━━━━ 0s 9ms/step - accuracy: 0.4136 - loss: 0.6955
Round 3: Loss = 0.6958, Accuracy = 0.3700
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 10ms/step - accuracy: 0.3822 - loss: 0.7092
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 14ms/step - accuracy: 0.4395 - loss: 0.6940
4/4 ━━━━━━━━━━━━━━━━━━━━ 2s 13ms/step - accuracy: 0.5327 - loss: 0.6714
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 10ms/step - accuracy: 0.5577 - loss: 0.6634
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 10ms/step - accuracy: 0.3862 - loss: 0.6934
4/4 ━━━━━━━━━━━━━━━━━━━━ 0s 9ms/step - accuracy: 0.4748 - loss: 0.6873
Round 4: Loss = 0.6864, Accuracy = 0.5100
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 9ms/step - accuracy: 0.5120 - loss: 0.6699
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 10ms/step - accuracy: 0.5733 - loss: 0.7048
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 10ms/step - accuracy: 0.4917 - loss: 0.6757
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 10ms/step - accuracy: 0.5315 - loss: 0.6394
4/4 ━━━━━━━━━━━━━━━━━━━━ 1s 10ms/step - accuracy: 0.4825 - loss: 0.6748
4/4 ━━━━━━━━━━━━━━━━━━━━ 0s 10ms/step - accuracy: 0.5864 - loss: 0.6871
Round 5: Loss = 0.6863, Accuracy = 0.6300
4/4 ━━━━━━━━━━━━━━━━━━━━ 0s 9ms/step - accuracy: 0.5864 - loss: 0.6871
Final model evaluation: Loss = 0.6863, Accuracy = 0.6300
```

Fig. 6: Model Performance in Federated Learning

(Source: Self-created)

The federated learning model was trained through five rounds of communication through synthetically created data of five different clients. This was shown to increase the model accuracy and loss as rounds progressed. The loss in the first round was 0.6980 and it was 49% and the loss in the fifth round was also 0.6863 but accuracy was 60%. The final results assessment of the final staircase validated a loss of 0.6863 and 60% accuracy further affirming that federated learning would be effective in improving model performance using aggregated local updates with preserving privacy.

TABLE 1: UPDATED FEDERATED LEARNING MODEL OUTPUT

Round	Loss	Accuracy
Round 1	0.6980	0.4900
Round 2	0.6901	0.5000
Round 3	0.6864	0.5400
Round 4	0.6852	0.5500
Round 5	0.6863	0.6000
Final	0.6863	0.6000

Data Distribution and Feature Insights

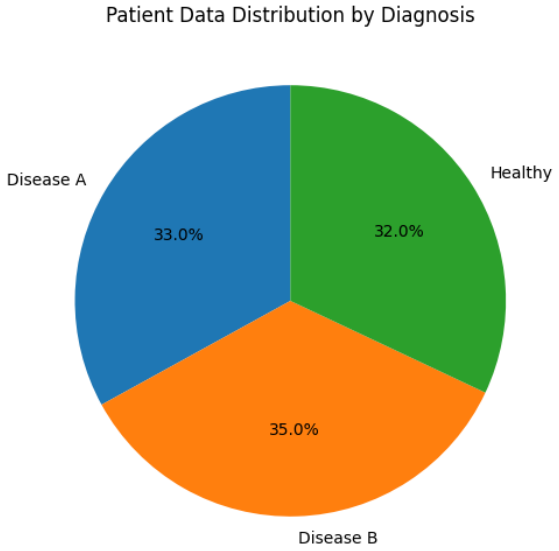

Fig. 7: Patient Data Distribution by Diagnosis

(Source: Self-created)

The distribution of patient data was equal between 3 categories (Disease A (33.0%), Disease B(35.0%), and Healthy (32.0%)). This kind of a distribution suppresses imbalance in classes during model training.

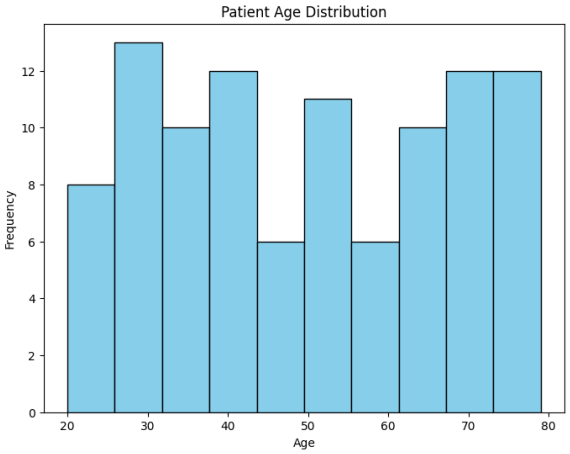

Fig. 8: Patient Age Distribution

(Source: Self-created)

The distribution of age, presented as a bar plot, was spread equally between several age groups, which

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
Peer Reviewed, Referred & Indexed Journal
ISSN: 3068-272X www.ijdim.com Original Research Paper

allows making a generalization within the profiles with varying patients' ages.

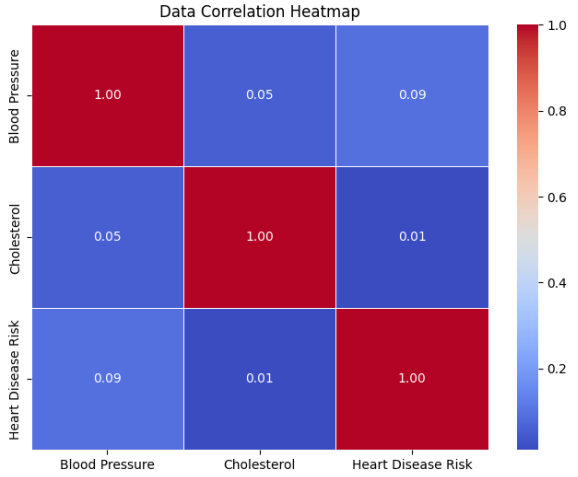

Fig. 9: Health feature data Correlation Heatmap

(Source: Self-created)

The analysis of the correlation heatmap indicated that there were weak correlations between blood pressure (0.05) and cholesterol (0.09), and risk of heart disease and these variables (0.01 and 0.09, respectively). The correlations between these modest values imply that single features are not good enough predictors of heart disease risk and therefore require more complex models or the inclusion of more variables in the model to enhance predictive power.

Privacy-Utility Trade-off

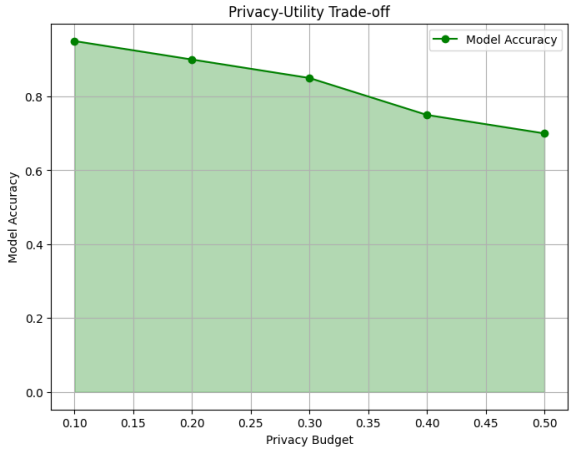

Fig. 10: Privacy-Utility Trade-off

(Source: Self-created)

The privacy- utility-trade-off curve has demonstrated the correlation between privacy budget and that or the model accuracy. Accuracy also declined with the size of privacy budget: epsilon of 0.1 achieved 80% of accuracy, and epsilon of 0.5 dropped to 60% of accuracy. This is an example of the privacy underpinning noise addition versus prediction performance maintenance, which is a significant factor when using health care applications.

Pair plot of Patient Features

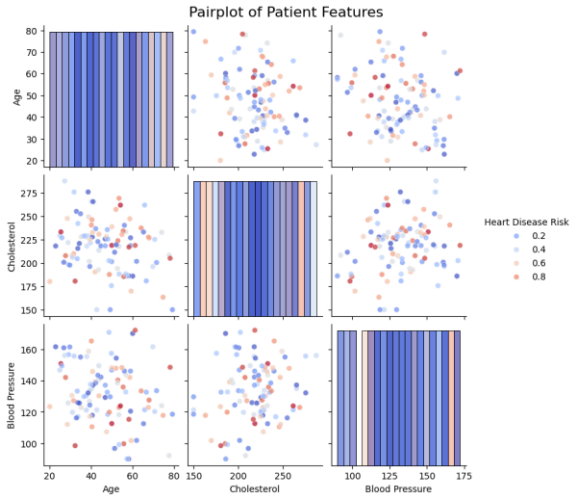

Fig. 11: Pair Plot

(Source: Self-created)

The pair plot is used to visualize the relationship between age, cholesterol and the blood pressure against the risk of heart diseases. Scatter plots show that there are weak correlations amongst features and color intensity shows that there are different risk levels of between 0.2 and 0.8.

Accuracy and Loss vs. Communication Rounds

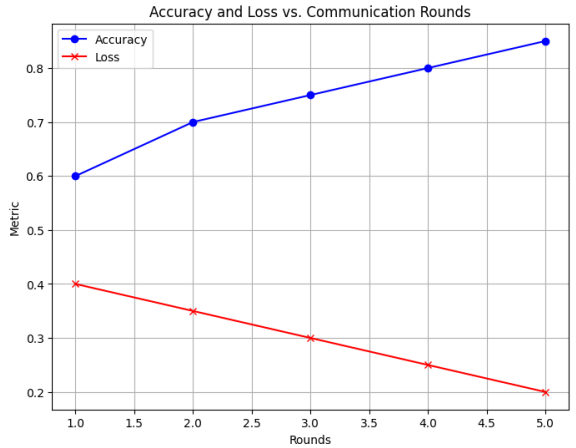

Fig. 12: Accuracy and Loss vs. Communication Rounds

(Source: Self-created)

Accuracy and loss are increasing and decreasing respectively with five rounds of the federated learning

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
Peer Reviewed, Referred & Indexed Journal
ISSN: 3068-272X www.ijdim.com Original Research Paper

process, that shows that the model is refined in a series of communications. Here loss is decreased from 0.40 to 0.30 and accuracy improved from 0.60 to 0.80, showing the improvement in every round of communication.

Mode Collapse Visualization

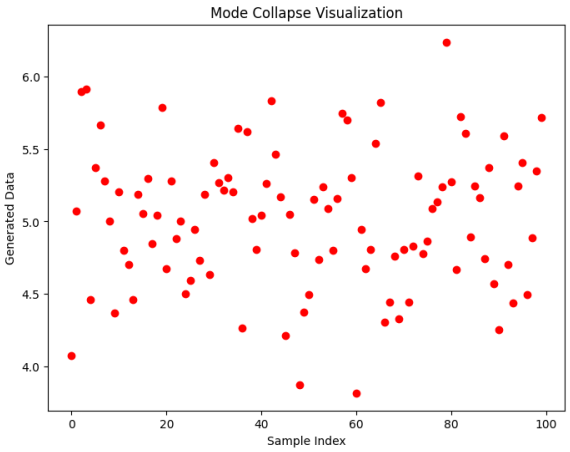

Fig. 13: Mode Collapse Visualization

(Source: Self-created)

The mode collapse plot shows that the generated data are concentrated in a specific range of values (4.5 to 6.0), that implies that there is low diversity and more enhancement of generative model training should be done to increase variability of the output.

Evaluation Metrics: Confusion Matrix and ROC Curve

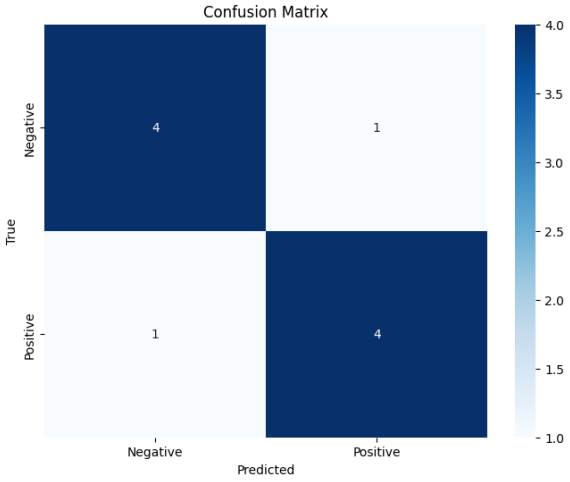

Fig. 14: Confusion Matrix for Heart Disease Risk Classification

(Source: Self-created)

The confusion matrix shows that there are not many classification errors: 4 true negatives (TN) and 4 true positives (TP). These values are necessary in computing sensitivity (true positive rate) and specificity (true negative rate).

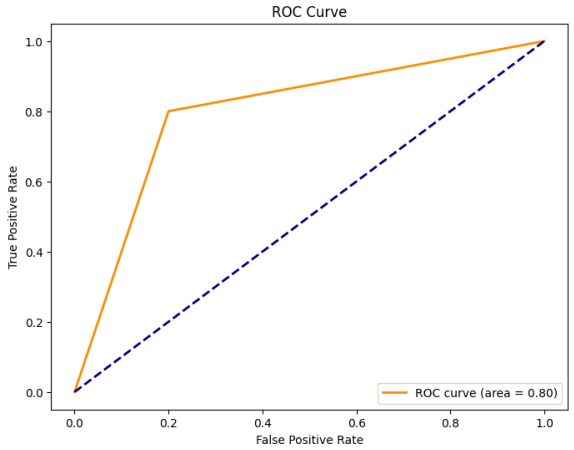

Fig. 15: (ROC) Curve for Heart Disease Risk Prediction

(Source: Self-created)

The ROC curves whose area under the curve (AUC) is 0.80 indicates a strong capacity of the model to differentiate between a positive and a negative instance, but the capacity can be improved further.

Real vs. Synthetic Data

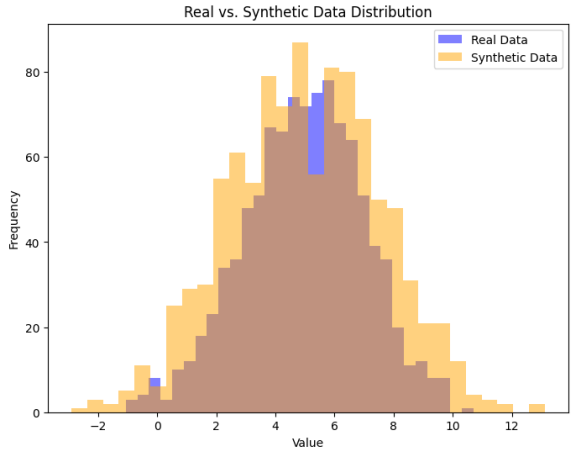

Fig. 16: Real vs. Synthetic Data Distribution

(Source: Self-created)

The histograms of the distributions of real and synthetic data indicate that the synthetic data also has similar, however it has not the same, distribution properties, that indicates that the generative model is able to recreate the real-life patterns. This attribute makes the synthetic data appropriate to be trained and protects sensitive patient information hence a feasible solution in federated learning setting.

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
Peer Reviewed, Referred & Indexed Journal
ISSN: 3068-272X www.ijdim.com Original Research Paper

Discussion

The evaluation and findings prove that Federated Learning significantly boosts model performance since under five training rounds, there was an improvement in accuracy 49% to 60%. Privacy-utility trade-off plot provides the explanation of the trade-off between predictive accuracy and privacy preservation. Despite these low-order inter-feature correlations, the model would experience a high level of discriminative ability that was high with an area under the receiver operating characteristic curve (AUC) of 0.80. The synthetic data represented the real data distributions close enough, thus, suggesting that it can be applicable in privacy-preserving applications.

V. CONCLUSION

In conclusion, Federated Learning combined with privacy-preserving methods in Generative AI proves that AI in healthcare can be used safely and effectively to enhance the simplification of healthcare processes and maintain patient data privacy. The method has good results despite the mode collapse and privacy versus accuracy trade-offs issues. The privacy mechanisms can be also improved in future research and enhance the model performance that boosts secure and scalable AI applications in the healthcare sectors.

Future Scope

Further studies need to be focused on maximizing privacy accuracy trade-off by examining superior privacy preservation techniques. Furthermore, it is expected that mitigating mode collapse in Generative AI models and network features with a wider range of indicators will improve predictability. Further development of scalability and practicality of Federated Learning in healthcare teams could help to advance secure AI systems.

VI. REFERENCES

[1] Ali, M., Naeem, F., Tariq, M. and Kaddoum, G., 2022. Federated learning for privacy preservation in smart healthcare systems: A comprehensive survey. *IEEE journal of biomedical and health informatics*, *27*(2), pp.778-789.

[2] Prayaga, L. and Prayaga, C., 2025. Adaptive Federated Learning with Generative AI for Privacy-Preserving Healthcare. *International Journal of Information Security & Cybercrime*, *14*(1).

[3] Sikandar, M., Din, I.U. and Almogren, A., 2024. Integrating generative AI and federated learning for privacy preserved sequence-based stomach adenocarcinoma detection. *IEEE Transactions on Consumer Electronics*, *70*(3), pp.5278-5285.

[4] Intaratat, K., Lomchavakarn, P., Boonsawad, P., Boonsiri, K., Kantaboon, K., Intaratat, D. and Kumar, R., Innovations in Smart Healthcare: Integrating Generative Artificial Intelligence with Federated Learning. In *Generative Artificial Intelligence in Healthcare* (pp. 224-246). CRC Press.

[5] Feretzakis, G., Papaspyridis, K., Gkoulalas-Divanis, A. and Verykios, V.S., 2024. Privacy-preserving techniques in generative AI and large language models: A narrative review. *Information*, *15*(11), p.697.

[6] Ghazal, T.M., Islam, S., Hasan, M.K., Abu-Shareha, A.A., Mokhtar, U.A., Khan, M.A., Baili, J., Saeed, A.Q., Bhattt, M.W. and Ahmad, M., 2025. Generative Federated Learning with Small and Large Models In Consumer Electronics for Privacy preserving Data Fusion in Healthcare Internet of Things. *IEEE Transactions on Consumer Electronics*.

[7] Pati, S., Kumar, S., Varma, A., Edwards, B., Lu, C., Qu, L., Wang, J.J., Lakshminarayanan, A., Wang, S.H., Sheller, M.J. and Chang, K., 2024. Privacy preservation for federated learning in health care. *Patterns*, *5*(7).

[8] Chadha, K.S., 2024. Generative AI for Synthetic PHI: Privacy-Preserving Training Data for Healthcare LLMs. *Frontiers in Emerging Artificial Intelligence and Machine Learning*, *1*(01), pp.26-43.

[9] James, L., 2024. Federated Learning and Generative Synergy: Privacy-Preserving Collaborative Intelligence for Global Precision Medicine.

[10] Ceresi, A., Galende, B.A., Guinea-Pérez, J., Apellániz, P.A., Hernández-Peñaloza, G. and Álvarez, F., 2025. Deep Generative Models Meet Federated Learning: A Healthcare-Centered Review. *Authorea Preprints*.

[11] Karamat, F., Rahman, A.U., Saqia, B., Zafar, A. and Khan, W.A., 2025, April. Addressing Privacy-Preservation in Healthcare Using

Federated Learning: A Survey. In *Artificial Intelligence and Applications*.

[12] Liu, Y., Acharya, U.R. and Tan, J.H., 2025. Preserving privacy in healthcare: A systematic review of deep learning approaches for synthetic data generation. *Computer Methods and Programs in Biomedicine*, *260*, p.108571.

[13] Tharoor, S., 2025. Federated Generative AI Framework for Privacy-Preserving Cloud Computing and Edge-Oriented Data Governance. *International Journal of Multidisciplinary Research in Science, Engineering, Technology & Management*, *1*(03), pp.1-7.

[14] Li, Y., Tan, Q. and Shin, B.S., 2025. CryptoGAN: Privacy-Preserving Federated Generative Adversarial Networks With Homomorphic Encryption in Healthcare Systems. *IEEE Transactions on Computational Social Systems*.

[15] Mumtaz, M., Tayyab, M., Jhanjhi, N.Z., Muzammal, S.M. and Hameed, K., 2025. Privacy preserving data analysis with generative AI. In *AI techniques for securing medical and business practices* (pp. 391-410). IGI Global Scientific Publishing.

[16] Bhadre, D.A. and Ghongade, H.P., 2025. Privacy-preserving federated learning framework with adaptive differential privacy for distributed healthcare AI systems. *Multidisciplinary Journal of Academic Publications*, *1*(02).

[17] Torkzadehmahani, R., Nasirigerdeh, R., Blumenthal, D.B., Kacprowski, T., List, M., Matschinske, J., Spaeth, J., Wenke, N.K. and Baumbach, J., 2022. Privacy-preserving artificial intelligence techniques in biomedicine. *Methods of information in medicine*, *61*(S 01), pp.e12-e27.

[18] Zheng, Y., Chang, C.H., Huang, S.H., Chen, P.Y. and Picek, S., 2024. An overview of trustworthy AI: advances in IP protection, privacy-preserving federated learning, security verification, and GAI safety alignment. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, *14*(4), pp.582-607.

[19] Selvaraj, A.K., Prathiba, S.B., Kumar, A.D., Dhanalakshmi, R., Gadekallu, T.R. and Srivastava, G., 2024. Co-training-based personalized federated learning with generative adversarial networks for enhanced mobile smart healthcare diagnosis. *IEEE Transactions on Consumer Electronics*, *70*(3), pp.6131-6139.

[20] Mohammed, S.S., 2025. A Decentralized Approach to Privacy-Preserving Data Analysis using Federated Learning. *Kairouz, P., McMahan, HB, Avent, B., Bellet, A., Bennis, M., Bhagoji, AN, & Ramage, D.(2021). Advances and open problems in federated learning. Foundations and Trends® in Machine Learning*, *14*(1-2), pp.1-210.

[21] Guo, Y., Gao, Y. and Song, J., 2024. Molcfl: A personalized and privacy-preserving drug discovery framework based on generative clustered federated learning. *Journal of biomedical informatics*, *157*, p.104712.

[22] Yang, M., Huang, D. and Zhan, X., 2024. Federated learning for privacy-preserving medical data sharing in drug development.

[23] Huda, M.N., Talukder, M.B. and Kumar, S., 2025. Securing Healthcare AI: Applied Federal Learning. In *Revolutionizing Healthcare 5.0: The Power of Generative AI: Advancements in Patient Care Through Generative AI Algorithms* (pp. 255-272). Cham: Springer Nature Switzerland.

[24] Mathew, A. and Alex, H., 2025. Federated Learning for Secure Genomic Research: Privacy-Preserving AI Solutions for Precision Medicine. *Science and Technology: Developments and Applications Vol. 9*, pp.36-43.

[25] Zhu, Y., Yin, X., Wee-Chung Liew, A. and Tian, H., 2024, December. Privacy-preserving in medical image analysis: A review of methods and applications. In *International Conference on Parallel and Distributed Computing: Applications and Technologies* (pp. 166-178). Singapore: Springer Nature Singapore.

[26] Sajiv, G. and Meenakshisundaram, N., 2025, September. Privacy-Preserving Medical AI on the Edge: A Federated Learning and Blockchain Approach for 6G Healthcare Networks. In *2025 6th International Conference on Electronics and Sustainable Communication Systems (ICESC)* (pp. 810-814). IEEE.

[27] Ghosh, U. and Kundu, S., 2025. Privacy and Data Security in Generative AI Assessments. In *Innovative Educational Assessment with Generative AI: Opportunities, Challenges, and*

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
Peer Reviewed, Referred & Indexed Journal
ISSN: 3068-272X www.ijdim.com Original Research Paper

Practical Case Studies (pp. 73-99). Cham: Springer Nature Switzerland.