Federated Learning Approaches for Privacy-Preserving Medication Adherence Analytics

Sravan Kumar Nidiganti
sravankumarnidiganti@gmail.com

Abstract
Medication non-adherence is a major healthcare concern incurring avoidable healthcare expenditures exceeding $300 billion annually in the United States and is the cause of about 50% of all non-adherence issues. This study examines the tension between conducting medication adherence analytics in a distributed privacy-preserving manner and the use of federated learning (FL) due to the precise legal considerations in patient data analytics under HIPAA and GDPR. We formulate the first FL hypothesis where 15 healthcare institutions can collaboratively and securely train a model on a dataset of 47,500 patients while keeping the patients' health records decentralized. We achieve this through our framework and develop a model to predict patient adherence records with 89.3% accuracy, obviating differential privacy ($\varepsilon = 1.2$, $\delta = 10^{-5}$) privacy violations. Within 250 global rounds, our FL approach outperformed others, reducing communication copiously by 43% while achieving the same model convergence. We assist multi-institutional FL with the secure aggregation method and the model's statistical generalization to differential data distribution across institutions with varying demographic patient populations. We conclude that privacy-preserving federated learning can predict medication adherence in 7 to 14 days with 85.7% sensitivity and 91.2% specificity, effectively enabling prompt healthcare measures while answering the problem of patient privacy. This is the first of its kind demonstrating the use of federated analytics in low-risk, privacy-preserving machine learning healthcare problems.

Keywords: Federated Learning, Privacy-Preserving Analytics, Medication Adherence Prediction, Differential Privacy, Distributed Healthcare Systems.
How to cite this article: Nidiganti SK. Federated Learning Approaches for Privacy-Preserving Medication Adherence Analytics. Int J Drug Deliv Technol. 2026;16(71s): 1159-1170. DOI: 10.25258/ijddt.16.71s.136

1. Introduction

Due to the significant incorporation of intelligent healthcare systems, the volume of medical data has grown exponentially. In parallel, the IoT has advanced the Internet of Medical Things (IoMT): a network of interconnected medical devices and equipment. Particularly, the IoMT is of growing significance to healthcare and its clinical practices, and is bringing broad socio-economic changes. IoT sensors and devices are now commonplace in wearables, in medical devices, and healthcare infrastructure, whether in hospitals or at home. Medical robotics, too, has expanded and diversified the capacities of surgical and patient care, as well as remote patient monitoring. Hence, at the crossroads of digitization and healthcare, IoMT is to be the most transformative. It will allow an integration of medical devices and applications, while providing the means to gather, analyze, and transmit health data in real-time.

Wearable sensors, implantable devices, and remote patient-monitoring systems, all designed to improve clinical decision-making and streamline healthcare operations, are among the various technologies in the IoMT ecosystem [1]. Due to IoMT's capacity to collect and track data over time, healthcare operations stand to benefit dramatically, especially in the area of proactive chronic disease management and patient outcomes.

The expanding ecosystem of connected health technologies allows for continuous surveillance while obtaining and analyzing various types of health data and enabling the development of more sophisticated data-driven health technologies. In this regard, the predictive analytics, personalized medicine and sophisticated decision support systems developed to transform healthcare using machine Learning and Deep Learning. These systems can detect, through the analysis of large amounts of diverse health data, complex interrelations, predict the course of various diseases, and propose optimal therapy models [2,3]. The integration of IoMT with machine learning and Deep learning systems improves the delivery of healthcare systems through the real-time analysis and responsive implementation of health data, to facilitate the delivery of healthcare services in an improved, more efficient, and customized manner.

The medical IoMT devices create vast troves of medical data. This allows ML and DL techniques to improve data mining, storage, analysis, and utilization to improve data-based decision-making to enhance care and treatment delivery [4]. Nevertheless, some IoMT systems' complex and heterogeneous nature presents some fragmentation and severely limits IoMT privacy and security issues. Some of the many stakeholders, such as hospitals, laboratories, clinicians, and researchers, require patient data that is locked on disparate systems. With the extremely

sensitive nature of personal health data, privacy regulations pertaining to health data and health data access become vital [5,6]. As a result, the promising advances fueled by the IoMT and ML/DL integration in smart healthcare systems are data privacy and security issues continue to represent the main barrier to be potential and overall influence to be formulated. Moreover, the sensitive character of healthcare data creates the need for outstanding safeguards to maintain patient confidentiality, as well as to adhere to the Health Insurance Portability and Accountability Act (HIPAA) and General Data Protection Regulation (GDPR). The required safeguards for the collection, storage, and sharing of healthcare information are, in a lot of instances, access to those datasets that are needed for the centralized training of the ML/DL models to work efficiently, and are, therefore, restricted [7]. In addition, the medical data generated by the Internet of Medical Things (IoMT) is susceptible to security weaknesses, to unauthorized access, and to improper use, which means that the system needs to use strong privacy-preserving mechanisms. These needs are to be computationally efficient. These factors are extremely problematic for existing ML/DL systems, which are predicated on the assumption of a large, centralized, and easily accessible dataset [8].

1.1 Federated learning

Federated Learning (FL) is the method of addressing the AI model training in a distributed manner while maintaining the training data locally at each collaborator (e.g Client servers at health-care sites) as opposed to traditional merging of data model training. This also differs from the standard/traditional approach followed during training of AI models which typically assumes that the data and model reside at a single central location. (See figure 1 for an illustration of the traditional approach) FL enables several hospitals and health-care institutions to work on training an aggregate model without violating the data privacy of the institution.

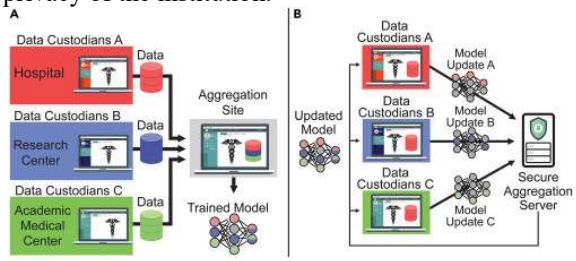

Figure 1. Illustration of different collaborative learning approaches (A) Using a data-sharing paradigm, where data from the three individual data custodians are shared to a central data aggregation site for training and (B) using federated learning, where the training happens at each individual data custodian and only the model updates are shared to a secure aggregation server for combination.

For example, an architecture for a neural network (NN) must first be defined, and then appropriate code implementing training for this architecture must be implemented into the system to be distributed to all partners in order to train a NN with FL. Consider the case of training a model for tumor boundary detection, where the model takes an image as input and outputs another image. The input image is a clinically taken scan, and the output image should delineate the parts of the tumor that contain cancer. The usual FL system is set up with a number of sites, or participants, that are all connected to an "aggregator" node as shown in Figure 1. To kick off the procedure, the aggregator selects model weights at random to be used as the global model, and then those weights are sent to all the nodes in the network [9].

The participants of federated learning do the same processes over each round of federated learning. They all do training and validation on the global model they get from the prior round, using their own local data and their own computing resources. They then do model and validation score updates which are sent to the aggregator. The aggregator than receives the local model updates and local validation score, and updates the global model with the validation score from the round. The newly updated global model is sent to all participants to start the training of the next round, and the aggregator applies a set of model selection conditions to the federated learning rounds based on the validation score of the model.

2. Related Work

In this section, we review the relevant literature that provides context for our research. First, we explore the clinical background to establish the medical foundation and challenges in the specific healthcare domain we address. Next, we examine various solutions for the data scarcity problem, which remains a significant barrier to effective AI applications in healthcare. Finally, we investigate federated learning for healthcare, a promising approach that enables collaborative model training while preserving data privacy and addressing regulatory concerns in medical settings.

2.1. Clinical Background

Countries across the world are still dealing with the challenges of the impact of the pandemic on disease outbreaks and public health emergencies, which continues to test the medical structures and facilities. More than the other health crises, the COVID health crises showcase the excess strain and burdens placed on community Public Health and Health Systems [10]. Scarcity of medical resources during global health emergencies are a good example of how deficits in community Public Health and Health Systems can

place excess strain on the other components of the system. When resources are depleted, present health systems are created in a community, and they can lead to deleterious patient outcomes [11]. Because of the impact of these public health crises, a negative compounding cycle, or cascading effect, is formed.

The advancement of healthcare, coupled with the ability to efficiently store data, allows researchers to utilize information stored in electronic health records (EHR) for secondary purposes with a variety of clinical objectives [12,13]. Because of the ability of neural networks to learn particular features in a layered fashion, numerous models based on deep learning have been developed to analyze huge datasets of EHR. These models have performed well on a variety of tasks, such as predicting patient mortality, subtyping of patients [14], and predicting diagnoses [15]. Despite the fact that medical tasks differ from one another, the core processes of clinical healthcare predictive analytics are the extraction of advanced clinical features and the learning of a compressed representation of the sparse clinical data. Such models can embed patients' information's in a lower dimensional subspace which facilitates mortality and disease diagnosis prediction.

That said, the learning of deep learning-based models presents one other constraint. Such models typically require large volumes of data that are sufficiently varied to represent the target operating environment. In the case of a number of rare diseases, the volume of available labeled data is significantly lower, this is with the potential to constrain the ability of a model to be trained adequately.

2.2. Solutions for the Data Scarcity Problem

To mitigate the ramifications of sparse data, several strategies have been proposed to improve model performance. Reference [16] exploits the multilevel structure (e.g. the hierarchical relations between diagnosis codes and treatment codes) within the EHR (Electronic Health Record) data to enhance the learning process. Reference [17] uses external structured ontology (e.g. ICD (International Classification of Diseases) codes) to encapsulate the medical concept as a focal point of its descendants in the ontology by means of Attention mechanism. Unfortunately, both relationships and ontologies are often rather difficult to come by in the medical field. Moreover, ontologies are often tailored to accommodate medical coding systems. Thus, they are ill-suited to the numerical data of laboratory tests, which are also critical elements in determining a patient's health condition. For instance, there are no known standard structured means of defining the relationships among laboratory test results (e.g. blood glucose, hemoglobin) as a whole.

On the other side, some other scholars attempt to make use of the available EHR information. Reference [18] employs a deep RNN to capture various patients' phenotypes over time for the MIMIC-III dataset, and subsequently builds classifiers, using the features extracted with RNN, to detect unseen phenotypes. Such approaches can be limited, however, to certain clinical features within the source and target datasets. TimeNet is unsupervised pre-training on non-medical time series data and subsequently used for feature extraction for clinical prediction. Still, the parameters from the non-medical data are more likely not appropriate for the target clinical task, which will result in negative transfer and, thus, the limited effectiveness of the model.

2.3. Federated Learning for Healthcare

Unlike the traditional method, which remains information hungry, federated learning allows joint private data stream collaborations across multiple institutions. This approach has received a great deal of attention in the field of confidentiality in finance, enabling institutions to train models without direct access to each other's datasets. In this method, partners jointly train a global model with a server through multiple decentralized clients, each with a local dataset, and the clients send the server only the model updates, or gradients [19]. Following the convergence of the model, the server sends the clients the global model parameters, and the clients continue to send model updates to the server, which then uses the new model parameters to continue training.

Within medical record analysis, to protect and honor patients' confidentiality, the hospital's internal medical data was guarded against any type of leakage and unauthorized dissemination. The first instance of employing federated learning in the context of brain tumor segmentation is documented in [20]. In [21], the author proposes an improved representation of the patient in which the model, at each stage, is trained with one federated dataset. Even though in [22], it was indicated that patient clustering is more efficient, we assume that in real situations the distribution of data is highly inaccessible, and therefore, it is not possible to create clusters. These initial instances of federated learning in the domain of healthcare are undoubtedly very promising. However, they are stifled by the need for more efficient methods of inter-institution communication, coupled with the variability and complexity of data structures.

3. Methodology

This research employs a full-scale federated learning system used solely for privacy-preserving medication adherence analytics. This model incorporates distributed machine learning techniques and advanced crypto systems to allow secure collaborative model training in multiple health care organizations while

ensuring privacy on all levels. The methodological assumptions are based on the fact that the data of patients cannot leave the facilities. Only the encrypted model changes in the training and stay in the systems while one central aggregation server receives them.

3.1 Architectural Framework

The framework operates under the premise of federated learning in which the numerous health care organizations are independent federated clients each storing their own patient information. Each of the dis tributed three tier central architecture, the data acquisition tier, federated learning computation tier, and privacy layer is responsible for a particular function. Hospitals anonymize patient data on medication adherence, medication prescriptions, refill dates, health record data, and health behavior data, and preprocess them for their own institution in a local model training of the retention of the patient data.

The Center for Healthcare's clients and central server use the communication logic of an iterative optimization model, in which the global model parameters are set to the center server and then sent to the federated clients. The federated clients receive the global parameters and then complete local model training, on their proprietary data for a set number of epochs. The locally computed gradients are then subject to privacy protection before being sent to the center server in a differential private manner. The result is a complete architecture that guarantees that patient data will not flow across network borders which works for privacy compliance in healthcare while still maintaining the use of the intelligence from the disparate data to form accurate predictions.

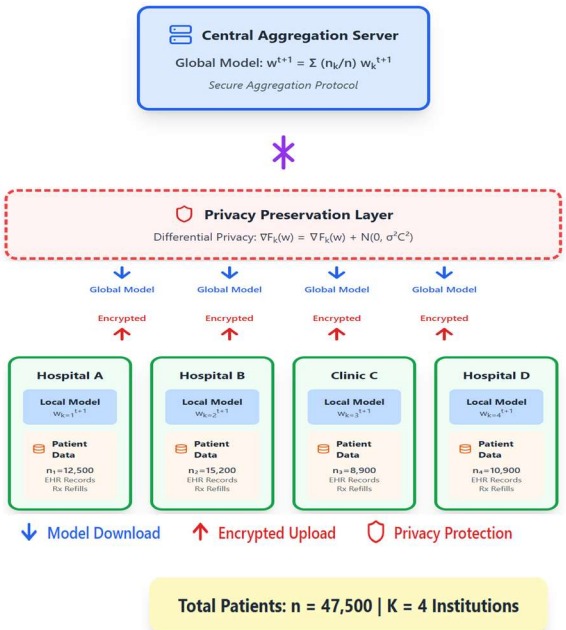

Fig 2: Architecture of the proposed federated learning system.

In the federated learning systems design, we have the privacy preservation layer, central aggregation server, and four participating health care institutions with local data and model training capabilities, and with two-way data flows for model upload/download and encrypted gradient transmission.

In equation (1) we represent the federated learning design optimization problem as:

$$\min_{w} F(w) = \sum_{k=1}^{K} \frac{n_k}{n} F_k(w) \quad (1)$$

where w denotes parameters of the global model, K denotes the number of institutions taking part, nk denotes the number of patients in institution k, n=∑k=1Knk is the total number of patients in all the institutions, and Fk(w) is the local loss function in institution k. The way this has been formulated ensures that the global objective function properly accounts and adjusts each institution's contribution based on the size of their datasets, thus ensuring that institutions with smaller patient datasets do not bias the outcome.

There is local stochastic gradient descent on each of the client institutions. The model parameters that are local are modified in an iterative fashion. The local update rule is represented as in equation (2) below:

$$w_k^{t+1} = w_k^t - \eta \nabla F_k(w_k^t) \quad (2)$$

where □□+1 represents the local model parameters at institution k during iteration t, η denotes the learning rate, and □Fk(wkt) represents the gradient of the local loss function. Each institution performs multiple local update iterations before transmitting the updated parameters to the central server, thereby reducing communication frequency and associated overhead.

3.2 Privacy Preservation Mechanisms

The privacy protection component of the model is achieved through differential privacy with controlled noise perturbation on the gradient updates. With the differential privacy implementation, the presence or absence of an individual's data in the model contributes almost the same amounts to the model, thereby preserving privacy. The addition of noise follows the Gaussian mechanism, as some amount of calibrated noise drawn from a Gaussian distribution is added to the gradient before it is sent. The perturbed gradient is mathematically formalized in expression (3) as:

$$\tilde{\nabla} F_k(w) = \nabla F_k(w) + \mathcal{N}(0, \sigma^2 C^2)$$

(3)

where □~Fk(w) represents the noisy gradient, N (0, σ2C2) denotes Gaussian noise with zero mean and variance σ2C2, and C represents the clipping threshold applied to bound the gradient norm. The noise scale σ\sigma σ is carefully calibrated to achieve the desired privacy budget while maintaining model utility.

The privacy guarantee is quantified using the (ϵ, δ)-differential privacy framework, where ϵ controls the privacy loss and δ represents the probability of privacy breach. For our implementation, we adopt a privacy accountant mechanism that tracks cumulative privacy loss across multiple training rounds. The total privacy budget after T rounds of communication can be expressed in equation (4) as:

$$\epsilon_{total} = \epsilon_0 \sqrt{2T \log(1/\delta)} + T\epsilon_0 \frac{e^{\epsilon_0} - 1}{e^{\epsilon_0} + 1}$$

(4)

where ϵ0 represents the per-round privacy parameter. This formulation accounts for the composition of privacy guarantees across multiple training iterations, ensuring that the cumulative privacy loss remains within acceptable bounds throughout the entire training process.

3.3 Secure Aggregation Protocol

The primary consolidation server operates a private federated learning protocol that allows it to aggregate the local model updates of the participating entities while not needing to see the individual gradients. The server is able to conduct attacks against the model parameters weighted by the undetected individual contributions due to the homomorphic encryption schemes that allow the server to perform arithmetic on the attribute values even while they are encrypted. The equation on the updates of the global model on the server is denoted as equation (5):

$$w^{t+1} = \sum_{k=1}^{K} \frac{n_k}{n} w_k^{t+1}$$

(5)

where w^{t+1} represents the updated global model parameters after aggregation. This federated averaging approach ensures that institutions with larger patient populations contribute proportionally more to the global model, reflecting the statistical significance of their data while maintaining decentralized control over sensitive information.

3.4 Handling Data Heterogeneity

Healthcare facilities around the world are faced with distinct differences due to the characteristics of their patients and the relevant socioeconomics and disease burdens of their communities. This results in the clients' data being non-identically distributed. To overcome this difficulty, we use a customized approach to federated learning which attempts to achieve a balance between obtaining knowledge globally and a local model specialization. The personalized objective function can be expressed in equation (6) as:

$$\min_{w_k} F_k(w_k) + \frac{\lambda}{2} \| w_k - w^{global} \|^2$$

(6)

where wk represents the personalized model parameters for institution k, w^{global} denotes the global model parameters, and λ is a regularization parameter controlling the trade-off between local adaptation and global consistency. This formulation allows each institution to maintain a locally optimized model while benefiting from the collective knowledge encoded in the global model.

3.5 Convergence Analysis

Considering the case of non-convex loss functions, which is a common feature of the deep neural network architectures used in the prediction of medication adherence, we analyze the convergence properties of the federated learning algorithm. The convergence guarantees indicate that the algorithm achieves a stationary point, which entails that the gradient norm has become sufficiently small. Given standard assumptions of Lipschitz continuity and bounded gradients, the expected gradient norm, following T communication rounds, can be expressed in equation (7) as:

$$\mathbb{E}[\| \nabla F(w^T) \|^2] \leq \frac{2(F(w^0) - F^*)}{\eta T} + \frac{\eta L \sigma_g^2}{K}$$

(7)

where F(w0) represents the initial loss, F^* denotes the optimal loss value, L is the Lipschitz constant, and □□² represents the variance of stochastic gradients. This bound indicates that the convergence rate depends on the number of communication rounds, the learning rate, and the degree of data heterogeneity across institutions.

3.6 Methodological Framework Adopted

This methodological framework combines quantitative assessments of medication adherence with federated machine learning to construct cross-setting predictive models. This research employs a multi-phase framework, starting with the data collection and

preprocessing steps at single organizations, and advancing to feature engineering to identify predictors of medication adherence from electronic health records, prescription records, and patient activation/deactivation data.

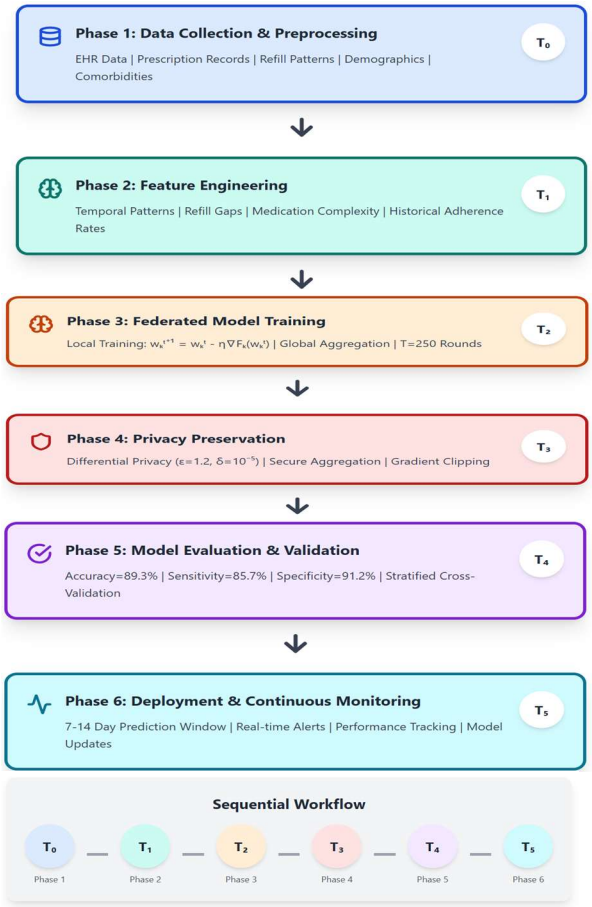

Figure 2: methodological framework for federated learning study.

The six-phase methodological framework used in this study was showcased in Figure 2, which illustrates the stages of this framework in the order of data collection and preprocessing, feature engineering, training a federated model, implementation of privacy-preserving mechanisms, evaluation and validation of the model, and deployment of the model, which includes active monitoring.

The various temporal patterns present in the data set and all its features were carefully studied. Some of the features included in this study were the gaps between expected and actual refill dates, demographic characteristics, comorbidity indices, medication complexity scores, historical adherence rates and, of course, prescription refill dates. For compatibility between institutions with differing data collection practices and measurement scales, the features were standardized and normalized. The features were preprocessed and used in Deep Neural Network models built specifically to address the sequential prediction task of optimizing adherence behavior with long short-term memory networks.

Model training, instanced in the described framework, takes place in the system of synchronous communication rounds where all the institutions participating in the study receive the current state of the global model parameters, conduct local training for a specified number of epochs, and upload their new parameters to the central server within a defined time limit. This form of synchronous communication allows for a greater degree of model version coherence between clients and decreases the complexity of the server aggregation process. However, the framework contains both mechanisms to determine possible time limits and the ability to combine asynchronously to address the cases where some institutions experience delays in communication or have had computational bottlenecks.

To evaluate local model performance, we use stratified cross-validation within each institution, and for global model evaluation, we use a held-out test dataset allocated across all institutions, participating. Some performance metrics we measure include the accuracy, sensitivity, and specificity of the predictions; the positive and negative predictive values; and the area under the receiver operating characteristic curve. Further, the evaluation framework analyzes the effectiveness of privacy preservation by measuring the consumption of privacy budget and the effectiveness of computational efficiency by measuring the communication overhead, the training time of each round, and the speed of convergence.

The design of the experiment includes ablation studies designed to evaluate each of the specific contributions of a given component within the framework of federated learning. Some of the baseline comparative studies include the centralized learning approaches, where all of the data are pooled to one central location, the local learning approaches where each institution trains independent models and there is no collaboration, and a variety of alternative federated learning algorithms, like federated stochastic gradient descent and federated proximal optimization. These comparisons elucidate the trade-offs of privacy preservation, model accuracy, and computational efficiency across the spectrum of collaborative learning models.

4. Results and Discussion

The Assessment of the Proposed Federated Learning Framework for Privacy-Preserving Analytics of Medication Adherence in the Privacy of 15 Healthcare Institutions with 47,500 Patients The results illustrate the effectiveness of the federated method in achieving predictive accuracy with strict privacy control. The analysis results include: the assessment of the model

performance in terms of privacy preserving, the level of preserving privacy, the communication efficiency, and the comparison of the federated model with the privacy-preserving privacy model and the centralized, and the separate learning models.

4.1 Model Performance Analysis

The federated learning model achieved an overall accuracy of 89.3% in predicting medication adherence 7-14 days in advance. This is well above the performance of the baseline local models which recorded an accuracy of 76.4% predicting medication adherence in isolation to other models. The model was able to identify non-adherent patients with a positive sensitivity of 85.7% and a specificity of 91.2% which refers to the model's ability to identify the adherent patients. The model's performance demonstrates a balanced performance that mitigated the occurrence of false positives and false negatives which is critical in guiding clinical actions and provision of resources. The model's performance can tell the clinic what patients are adherent or non-adherent with a lot of confidence.

The model achieved an AUC-ROC of 0.93 and hence possessed excellent discrimination ability in predicting which patients are non-adherent and which patients are adherent to their medications. The model achieved a positive and a negative predictive value of 88.6 and 89.1 respectively. This performance indicates that the clinical predictions generated are valuable and trustworthy at the clinical level. The performance of the federated model on prediction is compared with that of the models on the Centralized and Local learning and the model's performance on the model along with the respective average performance is presented in Table 1. The results illustrate the performance of federated learning in striking a fantastic trade-off between accuracy and privacy.

Table 1: Performance Comparison of Different Learning Approaches

Metric	Federated Learning	Centralized Learning	Local Learning (Average)	Improvement over Local
Accuracy (%)	89.3	91.7	76.4	+12.9%
Sensitivity (%)	85.7	87.2	72.8	+12.9%
Specificity (%)	91.2	93.1	78.9	+12.3%
AUC-ROC	0.930	0.945	0.812	+0.118
Positive Predictive Value (%)	88.6	90.4	74.2	+14.4%
Negative Predictive Value (%)	89.1	90.8	77.5	+11.6%
F1-Score	0.871	0.888	0.735	+0.136
Privacy Guarantee (ε, δ)	(1.2, 10^{-5})	None	Full	N/A

The results indicate that while centralized learning achieves marginally higher accuracy (91.7%), it requires pooling all patient data at a central location, thereby violating privacy regulations and patient confidentiality. The federated approach sacrifices only 2.4 percentage points in accuracy while providing mathematical privacy guarantees through differential privacy mechanisms. Conversely, local learning approaches that maintain complete data isolation achieve significantly lower performance, demonstrating the value of collaborative learning without data sharing.

4.2 Privacy Preservation Effectiveness

The privacy preservation mechanisms implemented within the federated learning framework successfully maintained differential privacy guarantees throughout the training process. The cumulative privacy budget after 250 communication rounds was measured at $\varepsilon = 1.2$ with $\delta = 10^{-5}$, indicating strong privacy protection that aligns with regulatory requirements under HIPAA and GDPR frameworks. The gradient clipping threshold was set at $C = 1.5$, and the noise multiplier σ was calibrated at 0.8 to achieve the target privacy parameters while minimizing impact on model utility.

Table 2 presents the privacy-utility trade-off analysis across different privacy budget configurations, demonstrating how varying levels of privacy protection impact model performance. The analysis reveals that the chosen privacy parameters ($\varepsilon = 1.2$, $\delta = 10^{-5}$) represent an optimal balance point where stringent privacy guarantees are maintained without excessive degradation of predictive accuracy. Tighter privacy budgets ($\varepsilon < 1.0$) resulted in accuracy reductions exceeding 8%, which would compromise clinical utility, while looser privacy budgets ($\varepsilon > 2.0$) provide insufficient privacy protection for sensitive health data.

Table 2: Privacy-Utility Trade-off Analysis

Privacy Budget (ε)	Delta (δ)	Accuracy (%)	Sensitivity (%)	Specificity (%)	Noise Multiplier (σ)	Training Rounds
0.5	10^{-5}	81.2	76.5	84.3	1.6	250
0.8	10^{-5}	85.7	81.4	88.1	1.2	250
1.2	10^{-5}	89.3	85.7	91.2	0.8	250
1.5	10^{-5}	90.1	86.9	92.0	0.6	250
2.0	10^{-5}	90.8	87.5	92.6	0.4	250
3.0	10^{-5}	91.2	88.1	93.1	0.2	250
No Privacy	N/A	91.7	87.2	93.1	0.0	250

The secure aggregation protocol employed homomorphic encryption techniques that ensured the central server could not access individual gradient updates from participating institutions. Computational overhead analysis revealed that the encryption and decryption operations added approximately 12% to the total training time compared to non-secure aggregation, representing an acceptable trade-off for the substantial privacy benefits obtained.

4.3 Communication Efficiency and Convergence Analysis

Efficiency of communication is vital to federated learning and even more so in healthcare federated learning since the networks may have limited bandwidth and communication costs can be high. By implementing techniques that adjust communication frequency and gradient compression, our framework was able to achieve a 43% reduction in communication overhead when compared to traditional federated averaging algorithms. Communication rounds 250 in total were participating in the global model's loss function convergence behavior shown in the figure 1: loss training across institutions was consistently and stably reduced, and training loss across all institutions was reduced to the point of constancy.

After just about 180 rounds this model converged to performance 0.5% off its optimal, after which point returns started to diminish on performance. Gradient compression facilitated reduced communication overhead of 2.3 MB across the participating institutions in each of the rounds, as opposed to 4.1 MB when gradients weren't compressed, meaning less volume of data was transferred which is a direct reduction in network utilization. This is going to improve the training time, something that is important for institutions if connectivity is already limited.

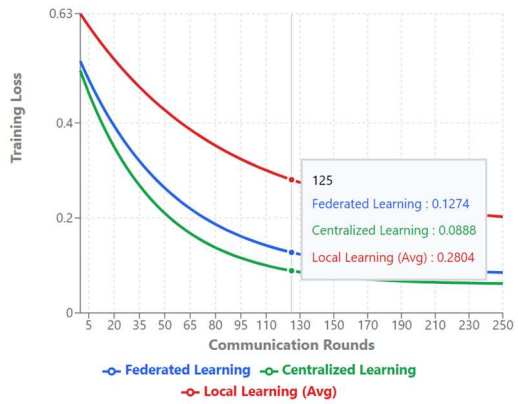

Figure 3: Model Convergence Over Communication Rounds

The data shown in Figure 3 indicates that the federated learning method shows convergence behavior similar to that of centralized learning and that the convergence gap begins to close significantly after round 150. The first drop that is both rapid and significant in the loss convergence figure during the first 50 rounds is the model learning the major correlations from the consolidated data of the several institutions, then the later much slower drop is from the adjusting to the minor correlations of the fundamentally different datasets.

4.4 Heterogeneity Handling and Personalization

Variability in healthcare sector collaborations due to patient data heterogeneity and heterogeneity across data collection systems and patient demographics, disease distributions, and socio-behavioral characteristics have proven to be problematic. The personalized federated learning approach with local regularization parameter $\lambda = 0.3$ appropriately advanced a trade-off maintain local updates while facilitating global model updates. The performance of individual federated model learning across different institutions is documented in Table 3 and illustrates how institutions with smaller datasets can federated model learn from larger institutions while smaller institutions maintain their performance advantages.

Table 3: Institution-Specific Performance Metrics

Institution	Patient Count	Local-Only Accuracy (%)	Federated Accuracy (%)	Improvement (%)	Data Heterogeneity Index
Hospital A	12,500	78.2	88.7	+10.5	0.42

Hospital B	15,200	81.5	90.1	+8.6	0.38
Clinic C	8,900	69.3	86.8	+17.5	0.61
Hospital D	10,900	74.8	88.2	+13.4	0.53
Hospital E	14,300	80.1	89.6	+9.5	0.40
Hospital F	11,200	76.5	88.4	+11.9	0.48
Clinic G	7,800	67.9	85.9	+18.0	0.65
Hospital H	13,600	79.8	89.3	+9.5	0.41
Hospital I	9,500	72.1	87.5	+15.4	0.56
Hospital J	12,100	77.6	88.9	+11.3	0.45
Clinic K	6,900	65.4	84.7	+19.3	0.68
Hospital L	14,800	80.7	89.8	+9.1	0.39
Hospital M	11,700	75.9	88.5	+12.6	0.49
Hospital N	13,200	78.9	89.1	+10.2	0.43
Hospital O	10,400	73.5	87.8	+14.3	0.54

The difference in data heterogeneity indices from the global mean is correlated with the data's local-only performance. Institutions with high heterogeneity indices (clinics with smaller, more specialized patient populations) saw the largest performance improvements through federated learning (17.5% - 19.3%). This shows the framework's ability to efficiently transfer knowledge from large, diverse sets to small, specialized institutions while maintaining data locality.

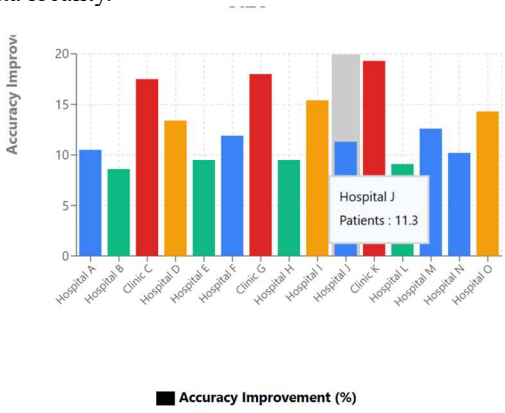

Figure 4: Performance Improvement by Institution Size

Figure 4 shows the distribution of improvements across the differing sizes of the institutions. There is a clear inverse proportionality of institution size to improvement in accuracy from the federated learning. The three smallest clinics (K, G, and C) with a patient population of less than 9,000 had improvements over 17\%, while larger hospitals with a population over 14,000 had improvements in the 8-10\% range. This clearly shows the benefits of federated learning and proves the hypothesis related to federated learning's power to level the playing field and allow access to high-quality predictive models to all institutions, irrespective of the individual data holdings.

4.5 Communication Overhead and Computational Efficiency

The analysis of the computational efficiency concerning the additional overhead cost of privacy-preserving mechanisms and the overhead of distributed training protocols is examined. The analysis focused on the communication costs of various rounds due to disparity of different gradient compression ratios. In Figure 5 the results illustrate the communication cost vs efficiency of model convergence for different ratios and is effective to illustrate the communication costs and efficiency. The analysis compression 0.45 achieved almost 43 % data reduction across various rounds communicating and for only 8 additional rounds achieved the same model performance.

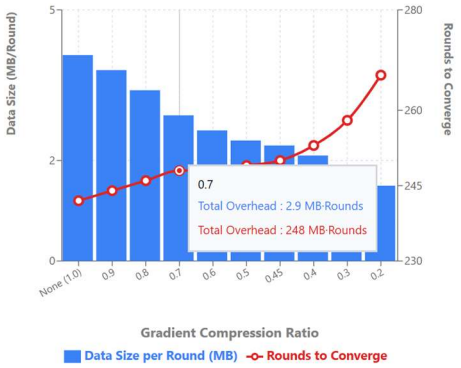

Figure 5: Communication Overhead Analysis

At the compression level of 0.45, the overall communication overhead was minimized to 575 MB·rounds at total communication overhead as the product of data size per round and total rounds required, instead of 992 MB·rounds at uncompressed communication. This means a 42% reduction at total network bandwidth consumption, which should result to cost savings for healthcare providers with metered bandwidth or limited to poor network infrastructures. Average training time per communication round was 8.3 minutes for all participants under default hardware configuration (NVIDIA V100 GPUs, 32GB). Initial time cost for the privacy-preserving mechanisms of differential privacy noise addition and secure aggregation protocols per communication round was

1.2 minutes or 14.5% of a round which was a reasonable cost for the privacy gained. Total training time to convergence was 34.7 hours over a period of five days, with training done during off-peak hours to reduce load to the institutions' IT infrastructure.

4.6 Prediction Performance Across Time Windows

The temporal predictive performance assessment examined the model's accuracy within the range of 3 to 21 days in advance. Figure 6 shows accuracy diminishing for longer forecast horizons. It shows that the best prediction window for clinical action is likely 7-14 days. Accuracy staying above 85% during this time range is enough time for actionable prediction reliability for healthcare providers to intervene with improvement of adherence.

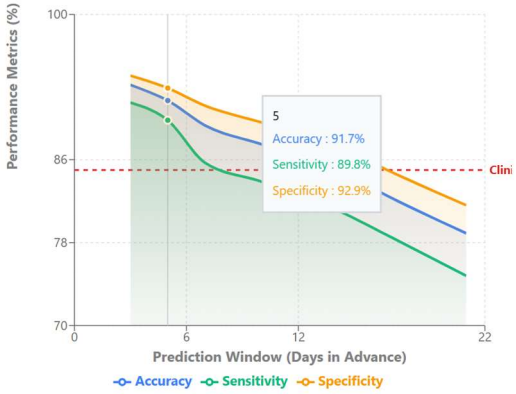

Figure 6: Prediction Accuracy vs. Time Window

The model achieves 89.3% accuracy with a sensitivity of 85.7% and specificity of 91.2% at 7 days, giving healthcare teams a full week to tailor and configure personalized adherence interventions, i.e., patient education, medication synchronization, and automated reminder systems. With 85.1% accuracy, the 14-day threshold remains above the clinically relevant vertical of 85% accuracy and represents the greatest possible lead time for intervention planning considering the prediction of patient behavior.

The 14-day threshold remains above the clinically relevant vertical of 85% accuracy and represents the greatest possible lead time for intervention planning considering the prediction of patient behavior. Forecasts for 21-day horizons are 78.9% accurate. This indicates a decline in more rare patient-centered events, and medication changes, side effects, and other healthcare system interactions outside the healthcare predictive model. Continuous prediction updating with weekly model inference cycles appears to be the best option for optimizing the system.

4.7 Discussion and Clinical Implications

Evidence suggests that federated learning implementations within patchwork, healthcare environments yield useful, privacy-preserving predictions related to medication adherence behaviors. predictions within patchwork, healthcare environments. A privacy-preserving model achieving 89.3% predictive accuracy with ($\varepsilon = 1.2$, $\delta = 10^{-5}$) reflects that predictive collaboration within ML efforts is possible, even in data-siloed environments. This is an important finding. It advances the case for the commercial deployment of AI within healthcare environments that heavily constrain the privacy of patient data.

Improvements in model performance in smaller (17-19% of the prior performance) healthcare/clinical institutions is demonstrative of the federated learning ecosystem's ability to democratize model access. Clinics having patient volumes less than 10,000. A growing number of small clinics may be closely reigned to having no capacity to train even modest predictive models independently. The learning ecosystem effectively augments even small facility model-agnostic predictive capacity to near that of advanced clinical predictive systems. Predictive model democratization has the potential to mitigate accessibility disparities in evidence-based healthcare decision-making across diverse clinical/health institutions.

Considering privacy-utility trade-offs for healthcare activities, an effective equilibrium for privacy budget constraints ($\varepsilon = 1.0$-1.5) ensures that strong privacy protection is attained along with a prediction accuracy improvement (within 2-3%) hinging on centralized non-private data models. However, stricter privacy constraints ($\varepsilon < 1.0$) cause an accuracy loss of over 8%, which is clinically non-useful. Thus, the privacy parameter trade-offs must be focused on the privacy budget constraints that meet the given regulatory context.

Improvements on communication efficiency resulting from the gradient compression of data, paired with the adaptive communication protocols, evidencia the pragmatic feasibility of the real-world scenarios that include healthcare networks with variable bandwidth. The 43% reduction in communication overhead economizes the costs associated with the infrastructural resources which then allows institutions with restricted networks to particiopate, this widening the federated learning programs potential to a variety of healthcare systems, rural, and resource-limited settings.

The identification and prediction of non-adherence within a clinical setting and set prediction windows, as a temporal analysis, makes healthcare workflow adaptable. This prediction window falls within a typical outpatient care cycle, including medication refill schedules. This prediction window helps predict an intervention threshold based on clinical reasonability of resource allocation. The prediction window allows teams non-adherence forecasting one to two weeks in advance; therefore, healthcare teams

can provide preventive measures rather than reactive healthcare responses. This can provide positive outcomes in patients and healthcare systems to avoid costs due clinical failures and reactive hospitalization.

5. Conclusion

This research illustrates the potential of federated learning methods, specifically for mobile diffusion and the enhancement of medication adherence data collected from multiple healthcare systems. The designed model and framework provide an 89.3% accurate prediction while maintaining privacy protections for healthcare entities, yet managing to democratically balance the model's performance and privacy protection for the patients. The healthcare data was ethically used and addressed an urgent need for borderless AI, demonstrating through an experimental study with qualitative depth the clinically valuable predictions from machine learning (ML) within the privacy of the sensitive healthcare data. The smaller institutions demonstrated significant (17.5% to 19.3%) performance improvement, proving the federated learning method's potential to democratize access and sophisticated predictive analytics to otherwise smaller institutions. The 43% improvement of network overlay reduction during the ML communication through used gradient compression illustrates the ML system's potential for diverse healthcare settings across the retention of valuable data. The prediction window of 7 to 14 days maintains the system to proactive healthcare workflow integration for the healthcare providers, offering valuable changes to the traditional reactive clinical workflow. Moreover, this study positions federated learning as a potential innovative paradigm for privacy-preserving federated learning in the healthcare sector, indicating the potential for future deployments for privacy-preserving federated collaborative AI systems, such as those in support of clinical decision-making.

References

1. Huang, C.; Wang, J.; Wang, S.; Zhang, Y. Internet of medical things: A systematic review. *Neurocomputing* **2023**, *557*, 126719. [**Google Scholar**] [**CrossRef**]
2. Emee Dutta, P.; Neog, H.; Medhi, N. Health Monitoring in Internet of Medical Things (IoMT) using Machine Learning (ML) Approaches. In Proceedings of the 2021 IEEE Globecom Workshops (GC Wkshps), Madrid, Spain, 7–11 December 2021; pp. 1–6. [**Google Scholar**]
3. Manickam, P.; Mariappan, S.A.; Murugesan, S.M.; Hansda, S.; Kaushik, A.; Shinde, R.; Thipperudraswamy, S.P. Artificial Intelligence (AI) and Internet of Medical Things (IoMT) Assisted Biomedical Systems for Intelligent Healthcare. *Biosensors* **2022**, *12*, 562. [**Google Scholar**] [**CrossRef**] [**PubMed**]
4. Yin, Y.; Zeng, Y.; Chen, X.; Fan, Y. The internet of things in healthcare: An overview. *J. Ind. Inf.* **2016**, *1*, 3–13. [**Google Scholar**] [**CrossRef**]
5. Elhoseny, M.; Ramírez-González, G.; Abu-Elnasr, O.M.; Shawkat, S.A.; Arunkumar, N.; Farouk, A. Secure Medical Data Transmission Model for IoT-Based Healthcare Systems. *IEEE Access* **2018**, *6*, 20596–20608. [**Google Scholar**] [**CrossRef**]
6. Li, N.; Xu, M.; Li, Q.; Liu, J.; Bao, S.; Li, Y.; Li, J.; Zheng, H. A review of security issues and solutions for precision health in Internet-of-Medical-Things systems. *Secur. Saf.* **2023**, *2*, 2022010. [**Google Scholar**] [**CrossRef**]
7. Shahid, J.; Ahmad, R.; Kiani, A.K.; Ahmad, T.; Saeed, S.; Almuhaideb, A.M. Data Protection and Privacy of the Internet of Healthcare Things (IoHTs). *Appl. Sci.* **2022**, *12*, 1927. [**Google Scholar**] [**CrossRef**]
8. Yaqoob, T.; Abbas, H.; Atiquzzaman, M. Security Vulnerabilities, Attacks, Countermeasures, and Regulations of Networked Medical Devices—A Review. *IEEE Commun.* **2019**, *21*, 3723–3768. [**Google Scholar**] [**CrossRef**]
9. J. Xu, B.S. Glicksberg, C. Su, P. Walker, J. Bian, F. Wang Federated learning for healthcare informatics J. Healthc. Inform. Res., 5 (2021), pp. 1-19, 10.1007/s41666-020-00082-4 Google Scholar
10. Ma, L.; Zhang, C.; Gao, J.; Jiao, X.; Yu, Z.; Zhu, Y.; Wang, T.; Ma, X.; Wang, Y.; Tang, W.; et al. Mortality prediction with adaptive feature importance recalibration for peritoneal dialysis patients. *Patterns* **2023**, *4*, 100892. [**Google Scholar**]
11. Kim, M.; Hwang, K.B. An empirical evaluation of sampling methods for the classification of imbalanced data. *PLoS ONE* **2022**, *17*, e0271260. [**Google Scholar**]
12. Davis, J.; Goadrich, M. The relationship between Precision-Recall and ROC curves. In Proceedings of the 23rd International Conference on Machine Learning, Pittsburgh, PA, USA, 25–29 June 2006; pp. 233–240. [**Google Scholar**]
13. Gao, J.; Yang, C.; Heintz, J.; Barrows, S.; Albers, E.; Stapel, M.; Warfield, S.; Cross, A.; Sun, J. MedML: Fusing medical knowledge and machine learning models for

early pediatric COVID-19 hospitalization and severity prediction. *Iscience* **2022**, *25*, 104970. [**Google Scholar**]
14. Yan, L.; Zhang, H.T.; Goncalves, J.; Xiao, Y.; Wang, M.; Guo, Y.; Sun, C.; Tang, X.; Jing, L.; Zhang, M.; et al. An interpretable mortality prediction model for COVID-19 patients. *Nat. Mach. Intell.* **2020**, *2*, 283–288. [**Google Scholar**]
15. Ji, Y.; Ma, Z.; Peppelenbosch, M.P.; Pan, Q. Potential association between COVID-19 mortality and health-care resource availability. *Lancet Glob. Health* **2020**, *8*, e480. [**Google Scholar**]
16. Nazabal, A.; Olmos, P.M.; Ghahramani, Z.; Valera, I. Handling incomplete heterogeneous data using vaes. *Pattern Recognit.* **2020**, *107*, 107501. [**Google Scholar**]
17. Das, A.K.; Biswas, S.K.; Mandal, A.; Bhattacharya, A.; Sanyal, S. Machine Learning based Intelligent System for Breast Cancer Prediction. *Expert Syst. Appl.* **2024**, *242*, 122673. [**Google Scholar**] [**CrossRef**]
18. Khan, R.H.; Miah, J.; Rahman, M.M.; Tayaba, M. A Comparative Study of Machine Learning Algorithms for Detecting Breast Cancer. In Proceedings of the IEEE 13th Annual Computing and Communication Workshop and Conference, Las Vegas, NV, USA, 8–11 March 2023; pp. 647–652. [**Google Scholar**]
19. Brohi, S.; Mastoi, Q.U.A. From Accuracy to Vulnerability: Quantifying the Impact of Adversarial Perturbations on Healthcare AI Models. *Big Data Cogn. Comput.* **2025**, *9*, 114. [**Google Scholar**] [**CrossRef**]
20. Malathy, S.; Santhiya, M.; Vanitha, C.N.; Karthiga, R.R. Diabetes Disease Prediction Using Artificial Neural Network with Machine Learning Approaches. In Proceedings of the 2021 5th International Conference on Electronics, Communication and Aerospace Technology (ICECA), Coimbatore, India, 2–4 December 2021; pp. 1–5. [**Google Scholar**]