Weighted Federated Learning with Encryption for Diabetes Classification

1st Puyang Zhao
Department of Biostatistics and Data Science
The University of Texas Health Science Center at Houston
Houston, USA
puyang.zhao@uth.tmc.edu

2nd Zhiyi Yue
Department of Biostatistics and Data Science
The University of Texas Health Science Center at Houston
Houston, USA
zhiyi.yue@uth.tmc.edu

3rd Xinhui Liu
Department of Statistics
The London School of Economics and Political Science
London, UK
x.liu146@lse.ac.uk

4th Jingjin Wu
Department of Statistics and Data Science
Beijing Normal-Hong Kong Baptist University
Zhuhai, P. R. China
jj.wu@ieee.org

***Abstract*—This study presents an innovative weighted Federated Learning (FL) framework with integrated encryption for diabetes classification across multiple healthcare institutions. Our comprehensive approach addresses three critical challenges in collaborative healthcare analytics: data privacy preservation, non-IID data distribution, and model performance optimization. The framework incorporates a weighted aggregation mechanism based on local data volumes to effectively handle client data imbalance, while implementing a lightweight masking-based encryption scheme to protect model parameters during transmission without compromising computational efficiency. We evaluate our approach using a comprehensive dataset of 15,347 entries from three internationally recognized medical organizations (ADCES, CDC, IDF) across five machine learning models: Logistic Regression (LR), Random Forest (RF), Support Vector Machine (SVM), three-layer Deep Neural Network (DNN), and deeper five-layer network (Deeper DNN). Experimental results demonstrate that weighted FL consistently equals or surpasses centralized learning performance while maintaining strict privacy compliance. Notable improvements include SVM AUC enhancement from 0.46 to 0.57 and RF AUC improvement from 0.70 to 0.76. The encryption mechanism introduces negligible overhead (0.0001s encryption, 0.0013s decryption per round) with minimal communication costs (0.16 KB per round). Our framework's ability to securely handle non-IID healthcare datasets while providing interpretable results through SHAP analysis positions it as a practical solution for privacy-preserving collaborative diagnostics. This research represents a significant advancement toward scalable, privacy-conscious medical analytics that can be adopted across diverse healthcare institutions without compromising data sovereignty or diagnostic accuracy.**

***Index Terms*—Encryption, Federated learning, centralized learning, Weighted aggregation, Diabetes classification**

I. Introduction

Diabetes, a global health crisis affecting over 530 million adults worldwide, demands innovative approaches to diagnosis and management that prioritize both accuracy and patient privacy [1], [2].While the abundance of healthcare data presents a significant opportunity for improving diagnostic precision, privacy concerns and regulatory restrictions often impede the aggregation and analysis of this valuable information across institutions. Traditional centralized approaches to data analysis face challenges in accessing diverse datasets due to these privacy constraints, potentially limiting the development of robust predictive models. In this context, FL emerges as a promising solution, offering a paradigm shift in how we approach collaborative model development in healthcare [3].

Recent studies have extended FL beyond conventional clinical applications, showing its adaptability in complex real-world scenarios. In industrial contexts, FL has been applied within 6G-enabled IIoT systems to support decentralized learning under energy and communication constraints, involving heterogeneous components such as smart devices and autonomous systems [4]. In healthcare, FL has enabled secure, collaborative training across distributed medical institutions by integrating with edge computing and IoT technologies, supporting low-latency decision-making while preserving data privacy [5]. These advancements highlight FL's growing relevance in privacy-sensitive, scalable applications.

FL enables the creation of robust machine learning models by leveraging diverse datasets from multiple institutions without compromising patient confidentiality, potentially enhancing predictive accuracy for conditions like diabetes. This approach addresses challenges faced by traditional methods, capturing a more comprehensive representation of the patient population while maintaining compliance with data protection regulation [6].

FL in healthcare faces challenges, including data heterogeneity, imbalance, and privacy risks. Our research addresses these issues through a Weighted FL framework that manages non-IID data distributions and incorporates encryption for enhanced privacy protection during model updates.

We propose an enhanced weighted FL framework that supports a broad spectrum of classifiers including LR, RF, SVM,

DNN, and Deeper DNN for the task of diabetes prediction. In the federated setting, local models are trained on-site and aggregated by a weight-based scheme that accounts for the relative data volume and class balance at each institution. An efficient encryption layer safeguards model parameters during transmission, thereby preserving patient privacy without exposing raw records.

This framework addresses three key challenges in collaborative healthcare analytics. First, it maintains privacy by keeping data within institutional firewalls and applying lightweight cryptography to model updates. Second, it improves data utility under non–IID conditions by weighting client contributions, which stabilizes training when local class distributions vary. Third, it provides architectural flexibility: both classical algorithms (LR, SVM, RF) and deep neural networks can be trained under the same protocol, easing deployment across sites with heterogeneous computational resources.

Our work extends prior studies that relied on centrally pooled datasets for diabetes classification with SVMs and RFs [7] and builds on recent FL initiatives focused on privacy in diabetes management [8]. By integrating weight-aware aggregation, support for deep learning, and an encryption mechanism with negligible overhead, we lay the groundwork for scalable, privacy-preserving diagnostic systems that can be adopted by multiple healthcare institutions without compromising data sovereignty.

The rest of this paper is organized as follows. Section II reviews related work on diabetes classification and FL. Section III details the weighted FL framework, covering data preparation, aggregation, and encryption. Section IV outlines the experimental setup and presents the results with interpretability analysis. Section V concludes and suggests avenues for future research.

II. Related Works

This chapter reviews relevant machine learning approaches for diabetes classification, with particular emphasis on FL frameworks and their applications in healthcare.

A. Machine Learning Approaches

Contemporary machine learning research employs three primary paradigms: centralized, distributed, and FL [9]. While centralized learning consolidates data in a single location for training, it faces significant challenges regarding data privacy and scalability [10]. Distributed learning addresses computational limitations through data partitioning across multiple nodes, but privacy concerns remain unresolved. FL emerges as a promising solution by enabling local model training on distributed devices, with only model updates shared through a central server [11]. This approach effectively balances computational efficiency with stringent privacy requirements essential for healthcare applications.

B. FL in Healthcare

FL encompasses various architectures, including horizontal, vertical, and federated transfer learning [3], [12]. The paradigm offers distinct advantages over traditional distributed approaches by accommodating data heterogeneity, varying node stability, and communication constraints while supporting privacy-sensitive environments through secure aggregation techniques [13], [14].

In diabetes management, existing FL research has shown promising but limited results. Boltri et al. [15] focus primarily on policy initiatives without addressing technological implementation challenges for cross-institutional data security. Islam et al. [8] develop a federated framework for diabetes complication prediction, emphasizing privacy and decentralization. However, their approach lacks advanced privacy mechanisms such as differential privacy and does not adequately address non-IID data distribution challenges common in healthcare settings.

C. Diabetes Classification Methods

Machine learning techniques have demonstrated effectiveness in diabetes classification using clinical datasets [7], [16]. Various algorithms have been evaluated, with SVMs showing reliable performance on the PIMA Indian diabetes dataset [7], while Bayesian networks achieved 99.51% accuracy in comparative studies [16]. Recent work by Butt et al. [17] and Maniruzzaman et al. [18] further validates the effectiveness of ensemble methods and neural networks, with Multilayer Perceptron achieving 86.08% accuracy across multiple classifiers including RF, LR, Naïve Bayes, and Decision Trees.

Despite these advances, existing approaches primarily focus on centralized learning paradigms, limiting their applicability in privacy-sensitive healthcare environments where data cannot be centrally aggregated. Furthermore, most studies do not address the inherent data heterogeneity and imbalance challenges present in multi-institutional healthcare settings.

D. Research Gap and Contribution

Building upon existing research, this study introduces a comprehensive FL framework that addresses key limitations in current diabetes classification approaches. Our contribution includes: (1) a weighted aggregation mechanism that handles non-IID data distribution across healthcare institutions, (2) an enhanced privacy protection scheme using masking techniques during model update transmission, and (3) improved model performance compared to traditional centralized methods while maintaining strict privacy compliance. This framework advances FL from basic analytics to a robust tool for collaborative diabetes management across multiple healthcare providers.

III. Methodology

This section presents our research methodology, detailing the data sources, preprocessing protocols, and the application of horizontal FL for diabetes classification.

A. Dataset Description and Pre-processing

We assemble a comprehensive diabetes dataset from three internationally recognised medical organisations: the Association of Diabetes Care & Education Specialists (ADCES),

the Centers for Disease Control and Prevention (CDC), and the International Diabetes Federation (IDF). These institutions were selected for their established credibility and capacity to provide high-quality, well-curated datasets. In total, our study encompasses 15,347 data entries, with the sample size from each organisation detailed in Table I. To address the

TABLE I
Dataset overview

Organisation	Sample size
ADCES	4 995
CDC	5 172
IDF	5 180

critical challenge of non-identical data distributions (non-IID) commonly encountered in FL environments, we developed a robust preprocessing protocol implemented locally at each participating node. This protocol encompasses: (1) imputation techniques to handle missing data while preserving data integrity; (2) normalisation procedures to standardise scales for effective aggregation and comparison across heterogeneous sources; and (3) rigorous anonymisation protocols to protect personal health information (PHI) in compliance with privacy regulations. By conducting these steps locally, we maintain data locality, minimise the risk of privacy breaches, and ensure sensitive information never leaves its originating organisation.

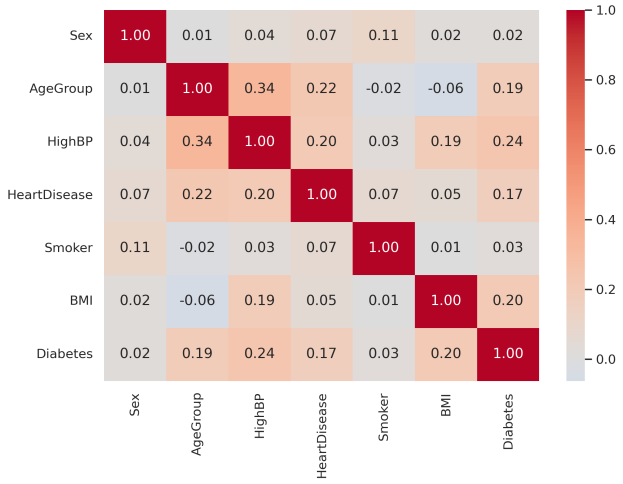

Fig. 1. Correlation matrix among variables including Diabetes status, BMI, Age group, HighBP, etc.

Figure 1 presents the Pearson correlation matrix for the seven clinical covariates considered in this study. Diabetes status exhibits its strongest linear association with hypertension (*HighBP*), yielding a coefficient of approximately 0.24. Moderate positive correlations are also observed with age group (≈ 0.19) and body mass index (BMI, ≈ 0.20). These findings indicate that older individuals, those with elevated blood pressure, and those with greater adiposity are disproportionately affected by diabetes in this cohort. By contrast, the correlation between sex and diabetes is negligible (≈ 0.02), and smoking shows only a weak yet positive relationship with diabetes. Collectively, the heat map highlights a coherent cardiometabolic risk cluster comprising age, hypertension, BMI, and diabetes, and it also confirms that none of the pairwise correlations exceed 0.40; therefore multicollinearity is unlikely to bias the subsequent multivariable modelling.

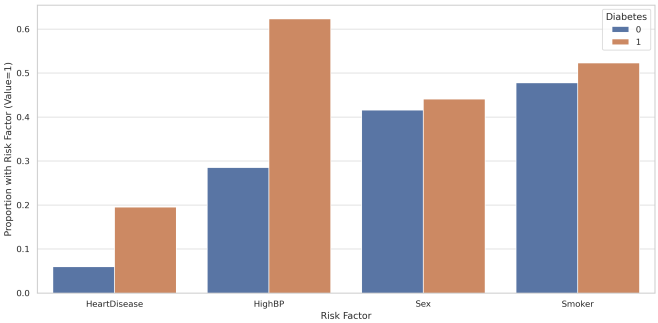

Fig. 2. Proportion of individuals with risk factors (e.g., high blood pressure, heart disease, smoking) in diabetic and non-diabetic groups.

Figure 2 compares the prevalence of four binary risk factors between individuals with and without diabetes. The bar plot indicates markedly higher rates of hypertension, heart disease, and smoking among participants diagnosed with diabetes. Hypertension is present in about 61 % of diabetic subjects compared with 29 % in the non-diabetic group, making it the most distinctive comorbidity. Heart disease appears in roughly 20 % of the diabetic cohort but only 7 % of their non-diabetic peers. Current smoking is reported by slightly more than half of people with diabetes, while fewer than half of those without diabetes smoke. In contrast, the proportion of men and women is almost identical in the two groups, suggesting that sex is unlikely to be a major confounder. These large prevalence gaps show that the dataset is strongly unbalanced for key cardiometabolic risk factors. Addressing this imbalance is essential to prevent biased inference in later analyses.

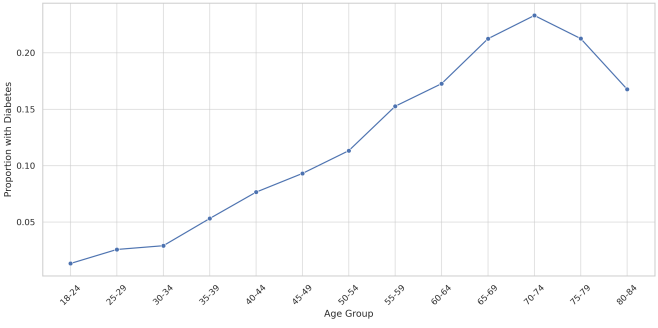

Fig. 3. The proportions of individuals with diabetes across age groups.

Figure 3 illustrates how the proportion of individuals diagnosed with diabetes changes across successive five-year age categories. Prevalence is minimal in early adulthood (18–34 years) and rises gradually through middle age. The increase

becomes steep from 55–59 years onward, culminating in the 70–74 year category where nearly one quarter of participants are affected. Although prevalence declines slightly beyond 75 years, it remains substantially higher than in younger cohorts. These data confirm that advancing age markedly elevates diabetes risk and establish age as one of the most important demographic drivers in this cohort.

B. Feature Selection and Experimental Setup

To prevent data leakage, feature selection is also performed within the federated setting. Each site independently computes Laplacian scores for its local data to identify the most informative features. The selected features include Sex, Age, HighBP, HeartDisease, Smoker, and BMI, with Diabetes as the binary target variable. Following feature selection, the dataset at each node is split into an 80% training set and a 20% test set, preparing it for the federated model training process.

C. Laplacian Score Computation

Our dataset's high dimensionality, comprising 14 distinct features, requires carefully selecting of the most informative attributes to enhance both computational efficiency and model accuracy. To achieve this, we utilize the Laplacian score method, which prioritizes features based on their capacity to preserve local neighborhood information. This method is especially effective for data with a manifold-like structure, as it maintains the inherent data topology within the selected feature subsets, thereby improving learning performance.

We construct a similarity graph where nodes represent samples and edges connect nearest neighbors based on feature similarity, represented by an adjacency matrix A. Let each sample $x_i \in \mathbb{R}^d$ be a feature vector in d-dimensional space, where $i = 1, \ldots, n$ and n is the total number of samples. From A, we derive a weight matrix S, with weights calculated as:

$$S_{ij} = \exp\left(-\frac{\|x_i - x_j\|^2}{t}\right), \tag{1}$$

where t is a scaling parameter.

The scaling parameter t is a pivotal hyperparameter that governs the exponential decay of affinity between data points. When t is small, the decay is steep and the weight matrix emphasizes only the nearest neighbors; when t is large, the decay is gradual and broader neighborhood information is retained. We determined t through an empirical cross-validation procedure on a held-out validation subset of the training data. A grid of candidate values was evaluated, Laplacian scores were computed for each candidate, and the value that maximized the variance of these scores was selected. This criterion increases the contrast between highly informative and less informative attributes, thereby producing the most discriminative feature subset for the downstream learning tasks.

Subsequently, we compute a diagonal matrix D where each diagonal element D_{ii} is the sum of the weights connected to node i, and define the unnormalized Laplacian matrix as:

$$L = D - S \tag{2}$$

For each feature f_r, where $f_r \in \mathbb{R}^n$ represents the values of the r-th feature across all n samples, the Laplacian score L_r is calculated to assess the relevance of the feature while accounting for data distribution. The score is computed as:

$$L_r = \frac{f_r^T L f_r}{f_r^T D f_r} \tag{3}$$

This score measures how effectively a feature preserves the local data structure, with lower scores indicating higher relevance. By systematically applying the Laplacian score method, we reduce the feature dimensionality while maintaining data integrity, thereby enhancing both our machine learning models' computational efficiency and predictive performance.

D. Comparison of Weighted FL and Centralized Learning

We aim to evaluate and compare the performance of machine learning models under two different learning frameworks: *Weighted FL* and *Centralized Learning*. The focus is on understanding the impact of data imbalance among clients and the benefits of weighted aggregation can bring in federated settings, particularly in handling non-IID (non-Independent and Identically Distributed) data distributions common in real-world scenarios.

1) Weighted FL:

a) Overview: FL is a decentralized machine learning paradigm where multiple clients collaboratively train a shared global model under the coordination of a central server while keeping the training data localized on each client to preserve privacy [19]. However, in practice, clients may have different amounts of data, leading to data imbalance, and the data distributions across clients may be heterogeneous. This heterogeneity poses challenges to the convergence and performance of the federated model.

To address these challenges, *Weighted FL* assigns different weights to each client's local model updates during the aggregation process proportionally to the number of data samples they possess. This weighting strategy ensures that clients with larger datasets have a more significant influence on the global model, potentially improving the overall performance and convergence stability.

b) Mathematical Formulation: Let us consider N clients, where each client i holds a local dataset D_i of size n_i. The total number of data samples across all clients is $n = \sum_{i=1}^{N} n_i$. At communication round t, each client performs local training initialized from the previous global model w^{t-1}, and computes a local model update Δw_i^t. That is, the updated local model is given by:

$$w_i^t = w^{t-1} + \Delta w_i^t, \tag{4}$$

where Δw_i^t denotes the local model update computed by Client i. The server then aggregates the weighted local updates to form the new global model:

$$w^t = w^{t-1} + \sum_{i=1}^{N} \frac{n_i}{n} \Delta w_i^t, \tag{5}$$

where $\frac{n_i}{n}$ reflects the contribution of client i in proportion to its local data volume. This update rule aligns the global model

with the underlying data distribution, potentially improving generalization.

c) Algorithmic Description: Algorithm 1 presents the detailed procedure of the Weighted FL algorithm.

Algorithm 1 Weighted FL Algorithm

Input: Local datasets $\{D_i\}_{i=1}^N$, number of communication rounds T, initial global model parameters w^0.
Output: Trained global model parameters w^T.
for each round $t = 1$ to T **do**
  for each client i **in parallel do**
    Local Training:
    Initialize local model parameters: $w_i^t \leftarrow w^{t-1}$.
    Train the local model on D_i to obtain updated parameters w_i^t.
    Compute local model update: $\Delta w_i^t = w_i^t - w^{t-1}$.
  end for
  Server Aggregation:
  Update global model parameters using weighted aggregation:

$$w^t = w^{t-1} + \sum_{i=1}^{N} \frac{n_i}{n} \Delta w_i^t.$$

end for
Return final global model parameters w^T.

d) Convergence Analysis: Weighted FL can improve convergence speed and model performance in settings with data imbalance. By giving more weight to clients with larger datasets, the aggregated global model can better reflect the overall data distribution. Previous works have shown that weighted aggregation can lead to convergence guarantees under certain conditions. For instance, Tian Li et al.'s study on the FedProx framework shows that adaptive computation based on system capabilities and effective management of statistical heterogeneity ensure convergence. [20].

e) Implementation Considerations: Implementing Weighted FL requires careful consideration of communication efficiency and computational overhead. The algorithm must ensure secure and efficient communication of model updates while handling potential issues such as stragglers and communication delays. Techniques such as compression of updates, asynchronous communication, and secure aggregation protocols can be employed to address these challenges.

2) Centralized Learning: In centralized learning, all client data from clients are collected and aggregated at a central server, where a global model is trained using the combined dataset. While this approach can potentially achieve high performance due to access to all data, it raises significant privacy concerns, particularly when dealing with sensitive medical records. The centralized model parameters w are obtained by minimizing a global loss function: Let $D = \bigcup_{i=1}^N D_i$ represent the aggregated dataset from all clients. The centralized model parameters w are obtained by minimizing a global loss function:

$$w^* = \arg\min_w \sum_{i=1}^{N} \sum_{(x_{ij}, y_{ij}) \in D_i} \ell(w; x_{ij}, y_{ij}), \tag{6}$$

where $\ell(w; x_{ij}, y_{ij})$ is the loss function for the sample (x_{ij}, y_{ij}).

3) Encryption Phase in FL: To enhance privacy, we incorporate an encryption phase using a masking technique during the transmission of local model updates. This method ensures that individual updates remain confidential.

a) Algorithm: Algorithm 2 details the encryption and aggregation process with masking.

Algorithm 2 Encrypted FL with Masking

Input: Local datasets $\{D_i\}_{i=1}^N$, number of rounds T, initial global model w^0, random seeds $\{s_i\}_{i=1}^N$.
Output: Trained global model w^T.
for each round $t = 1$ to T **do**
  for each client i in parallel **do**
    Local Training:
    Compute local model w_i^t based on D_i.
    Compute local update $\Delta w_i^t = w_i^t - w^{t-1}$.
    Encryption:
    Generate random mask m_i using seed s_i.
    Masked update: $\widetilde{\Delta w_i^t} = \Delta w_i^t + m_i$.
    Send $\widetilde{\Delta w_i^t}$ to the server.
  end for
  Server Aggregation:
  Aggregate masked updates: $\widetilde{\Delta w}^t = \sum_{i=1}^N \widetilde{\Delta w_i^t}$.
  Decryption:
  Compute total mask $M = \sum_{i=1}^N m_i$.
  Unmask aggregated update: $\Delta w^t = \widetilde{\Delta w}^t - M$.
  Update global model: $w^t = w^{t-1} + \frac{1}{N} \Delta w^t$.
end for
return w^T

b) Performance Metrics: We evaluate the encryption scheme based on the following metrics:

- **Average Encryption Time per Round:** Time taken to encrypt updates at each client.
- **Average Decryption Time per Round:** Time taken to decrypt aggregated updates at the server.
- **Average Communication Size per Round:** Total data transmitted between clients and server.

4) Centralized Learning: In centralized learning, data from all participating institutions are aggregated into a single dataset. This traditional method leverages the full spectrum of data to train machine learning models without considering data privacy or decentralization.

E. Machine Learning Models

We employ several machine learning models to evaluate the performance under both frameworks.

F. Machine Learning Models

We benchmarked five algorithms under both the centralized and federated frameworks.

LR: Implemented with `LogisticRegression`. In the federated setting, local coefficient vectors are averaged to form the global model.

RF: Implemented with `RandomForestClassifier`. For federated training, each site builds its own forest and the global prediction is the average of all local predictions.

SVM: Implemented with `SVC`. Similar to the random-forest procedure, aggregation in FL relies on averaging the decision scores from local models.

DNN: A three-layer feed-forward neural network built in `PyTorch`. Federated training follows the FedAvg protocol, where model weights are averaged after each communication round.

Deeper DNN: A five-layer neural network with a larger hidden dimension, also implemented in `PyTorch`. The same FedAvg weight-averaging strategy is used for global aggregation.

G. Experimental Setup

The dataset is split into training and testing sets, maintaining the class distribution in both sets. All features are standardized to improve model training performance.

For the FL framework, each client's local data is used to train a local model. The global model is updated by aggregating the local models using weighted averaging based on the number of samples at each client.

In the centralized learning framework, all data are combined to train a single model.

H. Communication Overhead Analysis

The communication overhead introduced by the encryption phase is analyzed in the FL framework. The communication size per round is a critical factor in assessing the efficiency of the FL system. Our analysis shows that the communication size remains relatively constant across training rounds, indicating a predictable communication overhead. This consistency in communication size benefits resource planning and system stability in FL implementations.

IV. Experimental Results and Analysis

Table II confirms that the weighted FL framework consistently equals or surpasses the centralized baseline, underscoring the benefit of collaborative training on distributed data. We first examine the classical models. The federated RF lifts AUC from 0.70 to 0.76 while maintaining a strong Weighted F1 of 0.82, indicating more reliable class separation. SVM shows an even larger gain: its AUC rises from 0.46 to 0.57 and its F1 improves from 0.79 to 0.81. LR likewise benefits, with higher precision and F1 under FL while retaining an AUC comparable to the centralized model.

Turning to deep learning, the pattern remains consistent. Federated training boosts the three-layer Deep Neural Network (DNN) from an F1 of 0.80 to 0.82 and keeps AUC at a competitive 0.77. The deeper five-layer network sees similar gains, confirming that the proposed weighting strategy scales effectively to high-capacity models.

Privacy preservation does not come at the cost of accuracy. Metrics for *Federated (With Encryption)* mirror those without encryption across every model. Moreover, the cryptographic layer is lightweight: encryption and decryption average 0.0001 s and 0.0013 s per round, respectively, with only 0.16 KB of communication—negligible overhead for routine clinical networks.

Because each institution retains its data locally, the framework remains robust under non-IID conditions. Local statistical properties are preserved while the global model benefits from diverse sources, thereby respecting privacy regulations and encouraging collaboration among sites with heterogeneous data.

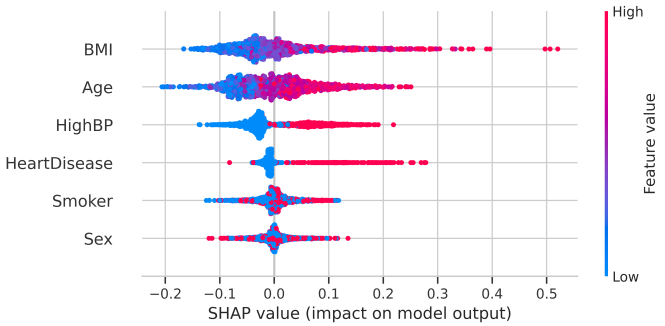

Fig. 4. SHAP Feature Impact for Diabetes Prediction.

To improve the transparency and interpretability of the prediction process, we applied the SHAP (SHapley Additive ExPlanations) framework to the best performing model, the RF classifier. Figure 4 displays the resulting SHAP summary plot and quantifies how each feature affects the predicted probability of diabetes. The horizontal axis shows the SHAP value: positive values raise the likelihood of a diabetes prediction, whereas negative values lower it. Each point represents a single observation and is colored by its original feature value.

The plot highlights Body Mass Index as the dominant driver: high BMI values align with large positive SHAP scores. Age is the second strongest contributor, followed by High Blood Pressure (HighBP) and a history of heart disease. In contrast, Smoking status and Sex cluster around zero, indicating limited influence for most cases. These effect patterns mirror established clinical knowledge and therefore lend biological plausibility to the model.

In summary, the weighted FL framework offers performance that meets or exceeds centralized learning, adds negligible encryption overhead, remains effective for both classical and deep architectures, and produces transparent, clinically meaningful explanations through SHAP. Together, these qualities position the approach as a practical and privacy-conscious solution for large-scale medical diagnostics.

TABLE II
Final performance comparison across learning frameworks

Model	Learning Framework	Weighted F1	AUC	Precision	Recall
LR	Centralized	0.81	0.78	0.81	0.86
LR	Federated (No Encryption)	0.82	0.76	0.82	0.87
LR	Federated (With Encryption)	0.82	0.76	0.82	0.87
SVM	Centralized	0.79	0.46	0.78	0.79
SVM	Federated (No Encryption)	0.81	0.57	0.89	0.87
SVM	Federated (With Encryption)	0.81	0.57	0.89	0.87
RF	Centralized	0.82	0.70	0.81	0.85
RF	Federated (No Encryption)	0.82	0.76	0.83	0.87
RF	Federated (With Encryption)	0.82	0.76	0.83	0.87
DNN	Centralized	0.80	0.78	0.83	0.86
DNN	Federated (No Encryption)	0.82	0.77	0.83	0.87
DNN	Federated (With Encryption)	0.82	0.77	0.83	0.87
Deeper DNN	Centralized	0.80	0.78	0.88	0.86
Deeper DNN	Federated (No Encryption)	0.81	0.77	0.76	0.87
Deeper DNN	Federated (With Encryption)	0.82	0.77	0.83	0.87

V. Conclusion

We developed a weighted FL framework with an integrated encryption phase for privacy preserving diabetes classification across multiple healthcare institutions. The framework supports a broad family of classifiers, including LR, RF, SVM, DNN, and Deeper DNN. Local training is performed on site. Model updates are aggregated with client weights that reflect the local sample size and class balance. A lightweight masking based encryption scheme protects model parameters during transmission so that raw patient records never leave the originating institution.

The study drew on curated data from three recognized organizations (ADCES, CDC, IDF) and applied a consistent local preprocessing pipeline that included imputation, normalization, and anonymization. Exploratory analyses identified a cardiometabolic risk cluster linking age, hypertension, body mass index, and diabetes, and showed that prevalence rises sharply after midlife. These data characteristics motivated our feature selection strategy and informed the interpretation of model outputs.

Comprehensive experiments demonstrated that weighted FL matches or improves upon centralized learning across most metrics for both classical and deep models. The largest discrimination gains appeared in the SVM and RF settings, where federated aggregation yielded stronger class separation under heterogeneous data. Deep neural networks trained federatively also performed competitively, indicating that the weighting scheme scales to higher capacity architectures. Performance remained stable when encryption was enabled. The measured encryption and decryption times were very small, and per round communication costs were minimal, suggesting the approach is feasible for routine clinical networks.

We further examined model transparency using SHAP explanations applied to the federated RF. Body mass index emerged as the leading driver of diabetes predictions, followed by age, high blood pressure, and prior heart disease. Smoking status and sex contributed little. These patterns align with established clinical knowledge and provide face validity for the learned decision rules. The ability to generate clinically grounded explanations is important for building trust in privacy preserving machine intelligence.

Several limitations merit attention. Our analyses used cross-sectional data and a limited feature set. Site level label quality and coding consistency may vary. Communication delays, client drop out, and stronger adversarial threat models were not fully explored. Although encryption overhead was small in our experiments, large-scale deployments may encounter bandwidth constraints and hardware heterogeneity. These issues define important directions for future work.

Future research should evaluate adaptive or personalized weighting schemes that respond to changing local data quality and class ratios. Stronger privacy protections such as differential privacy or secure multiparty aggregation could be layered onto the current protocol. Extending the framework to longitudinal outcomes, multimodal EHR data, and additional disease domains would test generalizability. Finally, prospective clinical studies that integrate workflow feedback and calibration monitoring are needed to assess real world impact.

In conclusion, this work shows that privacy aware collaboration across institutions is both practical and effective. Weighted FL with encryption can harness heterogeneous clinical data, protect patient confidentiality, produce interpretable models, and support scalable diagnostic decision making in healthcare analytics.

Acknowledgment

This work is partly supported by the Guangdong Provincial/Zhuhai Key Laboratory of Interdisciplinary Research and Application for Data Science, Project 2022B1212010006 and

in part by Guangdong Higher Education Upgrading Plan (2021-2025) UIC [R0400001-22] and [R0400024-22].

REFERENCES

[1] P. Zhao, X. Liu, Z. Yue, Q. Zhao, X. Liu, Y. Deng, and J. Wu, "Digan breakthrough: Advancing diabetic data analysis with innovative gan-based imbalance correction techniques," *Computer Methods and Programs in Biomedicine Update*, vol. 5, p. 100152, 2024.

[2] M. Lotfy, J. Adeghate, H. Kalasz, J. Singh, and E. Adeghate, "Chronic complications of diabetes mellitus: a mini review," *Current diabetes reviews*, vol. 13, no. 1, pp. 3–10, 2017.

[3] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *IEEE signal processing magazine*, vol. 37, no. 3, pp. 50–60, 2020.

[4] V. K. Quy, D. C. Nguyen, D. Van Anh, and N. M. Quy, "Federated learning for green and sustainable 6g iiot applications," *Internet of Things*, vol. 25, p. 101061, 2024.

[5] Q. V. Khanh, A. Chehri, A. D. Van, and Q. N. Minh, "Federated learning approach for collaborative and secure smart healthcare applications," *IEEE Transactions on Emerging Topics in Computing*, 2024.

[6] H. F. Ahmad, H. Mukhtar, H. Alaqail, M. Seliaman, and A. Alhumam, "Investigating health-related features and their impact on the prediction of diabetes using machine learning," *Applied Sciences*, vol. 11, no. 3, p. 1173, 2021.

[7] V. A. Kumari and R. Chitra, "Classification of diabetes disease using support vector machine," *International Journal of Engineering Research and Applications*, vol. 3, no. 2, pp. 1797–1801, 2013.

[8] H. Islam, A. Mosa *et al.*, "A federated mining approach on predicting diabetes-related complications: Demonstration using real-world clinical data," in *AMIA Annual Symposium Proceedings*, vol. 2021. American Medical Informatics Association, 2021, p. 556.

[9] A. M. Elbir, S. Coleri, and K. V. Mishra, "Hybrid federated and centralized learning," in *2021 29th European Signal Processing Conference (EUSIPCO)*. IEEE, 2021, pp. 1541–1545.

[10] M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauly, M. J. Franklin, S. Shenker, and I. Stoica, "Resilient distributed datasets: A {Fault-Tolerant} abstraction for {In-Memory} cluster computing," in *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*, 2012, pp. 15–28.

[11] L. Li, Y. Fan, M. Tse, and K.-Y. Lin, "A review of applications in federated learning," *Computers & Industrial Engineering*, vol. 149, p. 106854, 2020.

[12] C. Zhang, Y. Xie, H. Bai, B. Yu, W. Li, and Y. Gao, "A survey on federated learning," *Knowledge-Based Systems*, vol. 216, p. 106775, 2021.

[13] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings *et al.*, "Advances and open problems in federated learning," *Foundations and Trends® in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021.

[14] H. Ludwig and N. Baracaldo, "Introduction to federated learning," in *Federated Learning: A Comprehensive Overview of Methods and Applications*. Springer, 2022, pp. 1–23.

[15] J. M. Boltri, H. Tracer, D. Strogatz, S. Idzik, P. Schumacher, N. Fukagawa, E. Leake, C. Powell, D. Shell, S. Wu *et al.*, "The national clinical care commission report to congress: leveraging federal policies and programs to prevent diabetes in people with prediabetes," *Diabetes Care*, vol. 46, no. 2, pp. e39–e50, 2023.

[16] B. Alić, L. Gurbeta, and A. Badnjević, "Machine learning techniques for classification of diabetes and cardiovascular diseases," in *2017 6th mediterranean conference on embedded computing (MECO)*. IEEE, 2017, pp. 1–4.

[17] U. M. Butt, S. Letchmunan, M. Ali, F. H. Hassan, A. Baqir, H. H. R. Sherazi *et al.*, "Machine learning based diabetes classification and prediction for healthcare applications," *Journal of healthcare engineering*, vol. 2021, 2021.

[18] M. Maniruzzaman, M. J. Rahman, B. Ahammed, and M. M. Abedin, "Classification and prediction of diabetes disease using machine learning paradigm," *Health information science and systems*, vol. 8, pp. 1–14, 2020.

[19] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*. PMLR, 2017, pp. 1273–1282.

[20] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," *Proceedings of Machine learning and systems*, vol. 2, pp. 429–450, 2020.