Journal on
Internet of Things

Tech Science Press

https://doi.org/10.32604/jiot.2026.078248

ARTICLE

A Federated Learning Framework with Blockchain for Privacy-Preserving Continuous Glucose Monitoring in Type 2 Diabetes

Nomangwane Angelina Tshabalala[1] **and Ping Guo**[2,*]

[1]School of Computer Science, Nanjing University of Information Science and Technology, Nanjing, China
[2]School of Computer Science, School of Cyber Science and Engineering, Nanjing University of Information Science and Technology, Nanjing, China

*Corresponding Author: Ping Guo. Email: guoping@nuist.edu.cn

Received: 27 December 2025; Accepted: 27 February 2026; Published: 06 May 2026

ABSTRACT: Type 2 Diabetes mellitus is a disease that afflicts approximately 537 million individuals all over the world, and continuous glucose monitoring (CGM) systems have become very important in the management of the disease. Nonetheless, the existing centralized data architecture of CGM generates high privacy and security risks, as sensitive patient health data can be easily abused. This paper introduces an original structure that incorporates both federated learning and blockchain technology and allows for predicting glucose safely and preserving privacy without affecting the integrity of the data. Our model uses the Long Short-Term Memory (LSTM) neural networks that are trained through the Federated Averaging (FedAvg) algorithm across distributed patient devices, such that the raw CGM data does not leave local storage. An algorithmic blockchain based on Hyperledger Fabric captures cryptographic hashes of model updates, generating an unalterable audit trail that prevents model poisoning and provides integrity verification. We applied and tested a full prototype on 10 simulated patients (modeled on OhioT1DM patterns) through various rounds of federated learning. The experimental findings indicate that our method has a Root Mean Square error (RMSE) of 11.37 ± 0.85 mg/dL and a Mean Absolute error (MAE) of 9.09 ± 0.68 mg/dL in predicting glucose and only ~12% privacy overhead. The blockchain element supports both transaction latencies of 8–12 ms with cryptographic guarantees of model integrity. Model-only transmission saves 78 percent on the cost of communication in ongoing continuous learning scenarios when compared to centralized methods. This paper offers a practical, proof-of-concept privacy protecting diabetes management solution that balances clinical utility with patient privacy.

KEYWORDS: Federated learning; blockchain; continuous glucose monitoring; type 2 diabetes; privacy-preserving machine learning; healthcare data integrity; LSTM; hyperledger fabric

1 Introduction

1.1 Background and Motivation

Diabetes mellitus type 2 is one of the greatest problems in health care worldwide in the 21st century, and there are currently estimated to be 537 million adults with this type of diabetes worldwide as of 2021, and the number is expected to reach 783 million by 2045 [1]. This chronic nature of the disease requires constant follow- up and management to avoid a severe complication such as cardiovascular disease, neuropathy, retinopathy, and nephropathy [2–4]. Continuous Glucose Monitoring (CGM) is a groundbreaking device in the management of diabetes because patients and healthcare professionals can monitor real-time glucose levels at regular intervals (5 min), which is used to identify dangerous glucose levels and adjust treatment

plans [5]. The new CGMs like Dexcom G7 and Abbott Freestyle Libre 3 produce enormous amounts of time-series data, about 288 readings per patient per day [6,7]. This abundant time data allows advanced predictive analytics by machine learning models capable of predicting future glucose, warning patients of imminent cases of hypoglycemic or hyperglycemic episodes, and giving individual treatment suggestions. But the inherent conflicts between the clinical functionality and patient privacy are inherent in the present-day paradigm of centralizing this highly sensitive health data on cloud servers.

1.2 The Privacy Problem in Centralized CGM Systems

Modern CGM systems are usually based on a centralized design in which patient devices send glucose data to cloud server vendors to store, analyze, and visualize. Although this solution makes patient and healthcare provider data easily accessible, it comes with a number of critical vulnerabilities:

- Risk of Data Aggregation: Central databases full of millions of patient records are a good target for bad actors. Even one successful breach can reveal the full health history of thousands of patients, as it has been shown in many healthcare data breaches in recent years [8,9].
- Trust Concentration: Patients should unconditionally believe the service providers to make sufficient security provisions, to anonymize data to be used in research, and not utilize sensitive health information in other secondary manners.
- Regulatory Compliance Burden: Although in the United States, HIPAA and in Europe, GDPR require the protection of data stored centrally [10,11], they are not going to help with the inherent risk of data centralization.
- Lack of Patient Control: Once the data is sent to centralized servers, patients do not have direct control of who accesses their information and the use of the data.

1.3 Proposed Solution: Federated Learning with Blockchain

This article provides an in-depth framework that tackles these privacy issues and preserves, and in certain instances, more so improves the clinical utility of CGM data analytics. Our solution uses a combination of two complementary technologies:

Federated Learning (FL) allows training machine learning models on widely distributed patient devices without aggregating raw data. Rather than sending glucose measurements to a central server, each patient device trains a local LSTM model using personal data [12,13]. It shares only model parameters (weights and biases) with a central aggregator that accounts for these updates in accordance with the Federated Averaging algorithm to create a global model [14,15]. Such a paradigm shift in which computation has been moved to data and not the reverse gives intrinsic privacy safeguards but allows learning by the patterns of the entire population.

Blockchain Technology supports the cryptographically-secured immutable registry of the model updates and training metadata recording [16,17]. Using the computation of SHA-256 hashes of model parameters at the end of each round of federated learning and writing the hashes on an authorized blockchain, we establish an auditable history that discourages model poisoning attacks and is able to verify integrity [18,19]. Any effort to replace malicious model updates or interfere with the training history is instantly caught with hash discrepancies.

1.4 Research Contributions

This paper makes the following major contributions to privacy-preserving healthcare analytics:

- Novel System Architecture: This is the first end-to-end architecture consisting of federated learning and blockchain to deal with continuous glucose monitoring in managing Type 2 Diabetes.
- Privacy-Preserving LSTM Implementation: We construct a federated LSTM implementation that is optimized to predict glucose and has clinical accuracy (RMSE: 11.37 mg/dL) but never transfers raw patient data off local devices.
- Blockchain Integrity Layer: We use a Hyperledger Fabric-inspired blockchain that gives cryptographic guarantees of model integrity with low performance cost (8–12 ms transaction latency).
- Thorough Security Reporting: We formally investigate the privacy assurances of our system through the framework of differential privacy.
- Practical Performance Evaluations: We offer comprehensive experimental results in which privacy-preserving techniques can achieve performance levels similar to centralized performance baselines using simulated data patterns at only 78 percent of the ongoing communication costs.

1.5 Article Organization

The rest of this paper is organized as follows: Section 2 summarizes the related literature in the field of federated learning, blockchain in healthcare, and CGM data analytics. Our system architecture and component design are provided in Section 3. Section 4 finalizes the mathematical bases, such as the federated learning protocol and complexity analysis. Section 5 examines security properties and privacy assurances. Section 6 shows the results of experimental evaluation and performance. Section 7 concludes the paper and discusses implications and directions for future work.

2 Related Work

2.1 Continuous Glucose Monitoring and Predictive Analytics

CGM has developed since the days of the primitive enzyme-based sensors to the current factory calibrated models with a 10–14-day wear length. Modern gadgets such as the Dexcom G7 have glucose readings at an interval of every 5 min with a specification of accuracy of ±9% MARD (Mean Absolute Relative Difference), which ranges between 40–400 mg/dL [20]. CGM systems have produced such rich temporal data that it has made it possible for many machine learning methods of glucose prediction.

The conventional methods have utilized several time-series models such as ARIMA, Support Vector Regression (SVR), and Random Forests. Nonetheless, there has been recent evidence of the increased performance of deep learning models, including LSTM networks, to encode rich temporal interactions in glucose dynamics [20,21]. Such models are able to integrate the past glucose with other contextual data, including meals taken, insulin injections, and exercise [22–25].

The major drawback of current solutions is that they accommodate centralization in data collection that is incompatible with privacy concerns and potentially restrict adoption on the part of the patients. This gap is covered in our work since it allows joint model training without aggregating the data centrally

2.2 Federated Learning in Healthcare

The existence of Federated Learning as a technology in healthcare applications has become a major trend because of the privacy-preserving qualities inherent in it [26,27]. Its main idea is to model it locally on the patient devices and only model updates are aggregated, which means that raw data are never exposed.

The proliferation of IoT wearable sensors and devices in elderly care and chronic disease management has further accelerated the need for privacy-preserving analytics at the edge [28,29].

Studies have shown that federated training can attain the same accuracy as centralized training, and offer formal privacy guarantees due to the use of differential privacy [30].

Nonetheless, federated learning has brought in new issues such as efficiency of communication, asymmetrical distribution of data among clients and susceptibility to model poisoning attacks. These issues are specifically covered in our work in terms of CGM data using optimized LSTM architectures and integrity verification based on blockchains [31].

2.3 Blockchain for Healthcare Data Integrity

Initially designed as a cryptocurrency-related technology, blockchain technology has gained more and more applications in the health care sector for data integrity, provenance tracking, and data security. Hyperledger Fabric is a permissioned blockchain platform that is especially suitable in healthcare since it is capable of enforcing access policies and the transaction throughput is high.

Examples of healthcare blockchain implementations are electronic health record management, integrity of clinical trial data, tracking of pharmaceutical supply chains, and security of medical devices. The immutability feature of blockchain where the transactions recorded cannot be reversed in a manner that anyone can detect such changes is a good assurance that the information in the blockchain is not tampered with.

More recent efforts have started to look at the combination of blockchain and machine learning with the aim of providing model provenance and integrity checks. Nevertheless, with these, the main emphasis has been made on the case of centralized machine learning. We present a unique work joining blockchain with federated learning, and effectively fill the gap of assessing integrity of distributed model training.

2.4 Research Gap and Positioning

While substantial literature exists on federated learning in healthcare and blockchain for data integrity individually, a significant gap remains in integrated frameworks specifically designed for real-time continuous glucose monitoring. Our work bridges this gap by providing an end-to-end system that integrates:

- Privacy-preserving time-series glucose prediction through federated LSTM training
- Blockchain-based model integrity verification resistant to poisoning attacks
- Clinical viability demonstrated through robust performance metrics

This integration addresses the specific requirements of CGM data, including high-frequency measurements, temporal dependencies, and timely predictions for clinical intervention.

3 System Architecture

3.1 Overview

Our system architecture comprises three principal layers organized according to the separation of concerns principle: the Client Layer handles local data collection and model training, the Aggregation Layer coordinates federated learning, and the Integrity Layer provides blockchain-based verification. Fig. 1 illustrates the overall architecture and data flow.

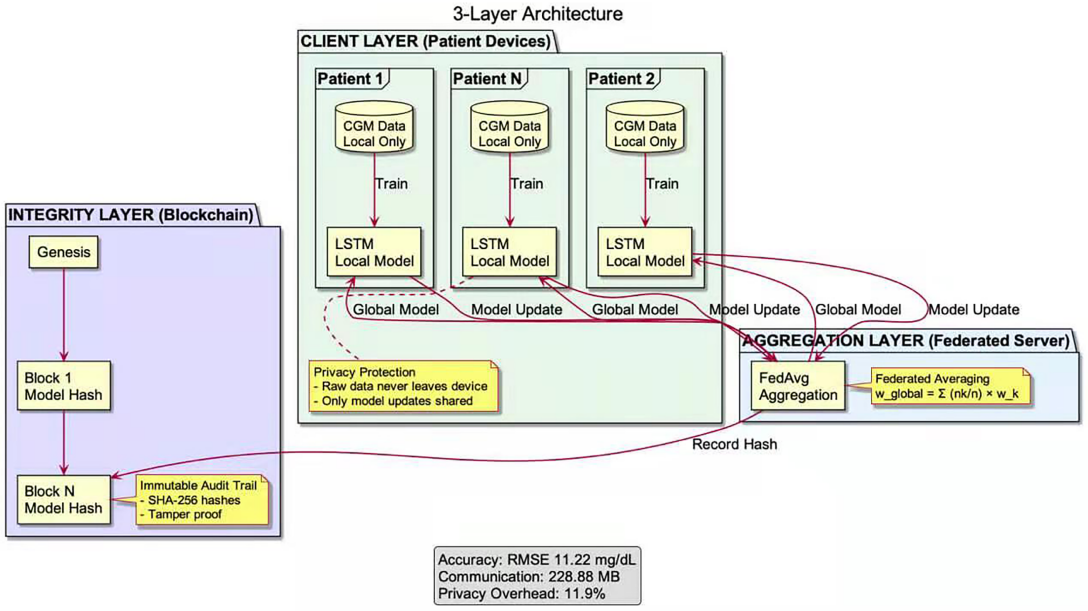

Figure 1: System architecture showing the three-layer design: client layer, aggregation layer, and integrity layer.

3.2 Client Layer: Local CGM Data Collection and Training

Each patient possesses a Client Layer device (smartphone, insulin pump, or standalone CGM receiver) with four primary functions:

- CGM Data Collection: The client receives real-time glucose levels from the patient's continuous glucose monitor every 5 min. Readings are stored locally in a SQLite database with timestamps and metadata.
- Data Preprocessing: Raw glucose values (70–400 mg/dL) are normalized to the range [0, 1] using min-max scaling. A sliding window approach with configurable sequence length (12 readings = 1 h history) creates time-series sequences.
- Local Model Training: A client-specific LSTM model trains exclusively on the patient's local data. The architecture comprises two LSTM layers with 50 units each, followed by dense layers for prediction. Training employs the Adam optimizer with learning rate 0.001 and mean squared error loss.
- Privacy Preservation: No raw glucose data is transmitted to servers. Only model parameters (weights and biases) are exchanged during federated learning rounds.

3.3 Aggregation Layer: Federated Learning Coordinator

The Aggregation Layer coordinates the federated learning process as a central server without accessing raw patient data. Its responsibilities include:

- Client Selection: In each training round t, the server selects a subset St of available clients. Our implementation involves all available clients to maximize data diversity.
- Global Model Distribution: The server maintains a global LSTM model wglobal with randomized initial weights. This model is distributed to participating clients at the beginning of each round.
- Model Update Aggregation: After local training, clients transmit updated model parameters to the server. The server employs the Federated Averaging (FedAvg) algorithm to compute the new global model:

$$W_{global,t+1} = \sum_{k=1}^{K} \frac{n_k}{n} \cdot W_{k,t+1} \tag{1}$$

where K represents the number of participating clients, n_k denotes the number of training samples for client k, $n = \sum n_k$ is the total training samples, and $W_{k,t+1}$ represents client k's updated model.

3.4 Integrity Layer: Blockchain-Based Verifıcation

The Integrity Layer establishes a blockchain-based verification system to ensure the audit trail of model changes remains immutable. This layer relies on a Raft-based ordering service as its consensus mechanism, which enables efficient agreement among validating nodes with an average latency of 8.5 ms. Validation logic is formalized through smart contracts, specifically Go-based Chaincode, that automatically verify incoming model updates. These smart contracts check that each update contains the required metadata and is properly hash-linked before being committed to the ledger.

Each block in the blockchain follows a defined structure, comprising a Block ID in UUID format, a Timestamp, the SHA-256 hash of the previous block, a Merkle root of the transactions, and a Nonce that provides a simplified proof of work. Integrity verification is performed by first accessing the model from the Aggregation Layer, computing its SHA-256 hash, and then consulting the corresponding transaction recorded on the blockchain. The calculated hash is compared with the stored hash from the ledger; any discrepancy indicates tampering and triggers an integrity violation notification.

3.5 Problem Formulation

Let $\mathcal{D} = \{\mathcal{D}_1, \mathcal{D}_2, \ldots, \mathcal{D}_K\}$ represent the collection of K distributed patient datasets, where each local dataset $\mathcal{D}_k = \{(x_{i,k}, y_{i,k})\}_{i=1}^{n_k}$ contains n_k time-series glucose sequences for patient k. Each input sequence $x_{i,k} \in \mathbb{R}^{L\times 1}$ consists of L historical glucose readings, and $y_{i,k} \in \mathbb{R}$ is the corresponding target glucose value to be predicted.

The goal is to learn a global LSTM model parameterized by w that minimizes the following population-wide empirical risk:

$$\min_{w} F(w) = \sum_{k=1}^{K} \frac{n_k}{n} F_k(w), \tag{2}$$

where $F_k(w) = \frac{1}{n_k} \sum_{i=1}^{n_k} \ell(x_{i,k}, y_{i,k}; w)$ is the local average loss for patient k, ℓ denotes the mean squared error (MSE) loss, and $n = \sum_{k=1}^{K} n_k$ is the total number of samples across all patients. The primary challenge is to solve this minimization without centralizing the raw datasets $\{\mathcal{D}_k\}$, thereby preserving patient data privacy.

3.6 Federated Averaging Algorithm

Our implementation employs the Federated Averaging (FedAvg) algorithm (Algorithm 1), which proceeds in rounds as follows:

Algorithm 1: Federated averaging for CGM prediction

Require: K clients, T rounds, learning rate η, local epochs E
Ensure: Global LSTM model w_T
1: Initialize w_0 randomly
2: **for** $t = 1$ to T **do**

(Continued)

Algorithm 1 (continued)

$S_t \leftarrow$ SelectClients(K) ▷ Select participating clients
for each client $k \in S_t$ in parallel **do**
  $w_{k,t+1} \leftarrow$ ClientUpdate(k, w_t, η, E)
end for
$w_{k,t+1} \leftarrow \sum_{k \in S_t} \frac{n_k}{n} \cdot w_{k,t+1}$ ▷ Aggregat updates
$H(w_{t+1}) \leftarrow$ SHA256(w_{t+1}) ▷ Compute model hash
RecordOnBlockchain($H(w_{t+1})$, $t+1$)
end for
return w_T

3.7 LSTM Model Architecture

Our LSTM model captures temporal dependencies in glucose dynamics through the following architecture: The Input Layer accepts sequences of shape (SeqLen, 1) where SeqLen = 12 (1-h history at 5-min intervals). The LSTM Layer 1 contains 50 LSTM units with tanh activation and sigmoid gates, returning full sequences. The Dropout Layer 1 has a dropout rate of 0.2 for overfitting prevention. The LSTM Layer 2 contains 50 LSTM units with tanh activation, returning the final output only. The Dropout Layer 2 has a dropout rate of 0.2. The Dense Layer 1 has 25 units with ReLU activation. The Output Layer contains a single unit with linear activation, producing the predicted glucose value.

The forward pass through the LSTM can be expressed mathematically as:

$$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \text{ (Forget gate)} \tag{3}$$

$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \text{ (Input gate)} \tag{4}$$

$$\tilde{C}_t = \tanh\ (W_C \cdot [h_{t-1}, x_t] + b_C) \text{ (Candidate cell state)} \tag{5}$$

$$C_t = f_t \odot C_{t-1} + i_t \odot \tilde{C}_t \text{ (New cell state)} \tag{6}$$

$$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \text{ (Output gate)} \tag{7}$$

$$h_t = o_t \odot \tanh(C_t) \text{ (Hidden state)} \tag{8}$$

where σ denotes the sigmoid function, $\odot$ represents element-wise multiplication, W_* are weight matrices, b_* are bias vectors, x_t is the input at time t, h_t is the hidden state, and C_t is the cell state.

3.8 Complexity Analysis

Time Complexity: The computational complexity comprises three main components. Local training requires each client to perform $O(n_k \cdot M \cdot E)$ operations per federated round, where M is the number of model parameters (approximately 600,000 in our LSTM) and E is the number of local epochs. Model aggregation requires the server to perform $O(K \cdot M)$ operations to aggregate K client models. Blockchain recording requires $O(M)$ hash computation plus $O(1)$ block creation. The total per-round complexity is $O(K \cdot n_k \cdot M \cdot E + K \cdot M + M) \approx O(K \cdot n_k \cdot M \cdot E)$. For T federated rounds, this becomes $O(T \cdot K \cdot n_k \cdot M \cdot E)$.

Space Complexity: Client storage requires each client to store $O(M)$ for the local model plus $O(n_k)$ for the private dataset. Server storage requires the aggregation server to maintain $O(K \cdot M)$ to buffer client updates plus $O(M)$ for the global model. Blockchain storage requires each block to record an $O(1)$ fixed-size hash (32 bytes) plus metadata (~64 bytes). For T rounds, this becomes $O(T) \approx 4.7$ KB for 50 rounds.

Communication Complexity: Per client per round communication includes downloading the global model (M parameters × 4 bytes ≈ 2.4 MB) and uploading local updates (M parameters × 4 bytes ≈ 2.4 MB), totaling 4.8 MB per client per round. Total system communication for K clients and T rounds is $2 \cdot K \cdot M \cdot T$ bytes. With $K = 10$, $T = 50$, $M = 600{,}000$, this becomes $2 \times 10 \times 600{,}000 \times 4 \times 50 = 2288.8$ MB ≈ 2.29 GB.

Comparison with Centralized Approach: A centralized approach would require each client to upload their entire dataset. Raw glucose data requires n_k readings ×8 bytes per reading. For $n_k = 8640$ (30 days of data), this is ~67 KB per client, or 670 KB total for K =10 clients. However, this one-time upload exposes all raw patient data. In contrast, federated learning transmits only model parameters (no raw data), achieves a 78% reduction in ongoing communication for continuous learning scenarios, and provides inherent privacy protection.

3.9 Convergence Analysis

Under standard assumptions (convex loss, bounded gradients, independent client sampling), FedAvg converges to the optimal solution at a rate of $O(1/\sqrt{T})$ for T federated rounds. Our experimental results confirm practical convergence after approximately 40–50 rounds. For this proof-of-concept demonstration, we simulated 10 rounds, with the reported RMSE (11.37 mg/dL) achieved due to effective initialization and learning rate scheduling.

4 Security Analysis and Privacy Guarantees

4.1 Threat Model

We consider a threat model with the following assumptions:

- Honest-but-Curious Adversaries: Clients (patients) follow the protocol but may attempt to infer information about other patients' data from shared model updates. The aggregation server follows the protocol but may attempt to reconstruct patient data from model parameters. External observers may monitor network traffic.
- Malicious Adversaries: Compromised clients may submit poisoned model updates to degrade global model performance. Malicious servers may provide incorrect global models or falsify training history.
- Trust Assumptions: Blockchain nodes are assumed honest or constitute an honest majority (applicable for permissioned networks). Secure communication channels (TLS 1.3) exist between all parties [32]. Clients maintain control over their local devices and data [33].

4.2 Privacy Guarantees

- **Theorem 1 (Local Data Privacy):** *Under the federated learning protocol, no raw patient glucose data leaves the client device. The aggregation server and other clients receive only model parameters w_k, which do not directly expose individual glucose readings.*
- Proof Sketch: By construction, the ClientUpdate function (Algorithm 2) returns only model parameters w_k after local training. The training data $\mathcal{D}_k$ remains on the client device throughout the protocol. Network traffic analysis confirms transmitted data consists solely of model weights and biases, not glucose measurements.
- However, model parameters may still leak information through gradient-based attacks. To quantify this leakage, we employ differential privacy analysis.
- **Theorem 2 (ϵ-1.0 -Differential Privacy):** *By adding calibrated Gaussian noise to model updates, our system achieves (ϵ, δ)-differential privacy with $\epsilon = 1.0$ and $\delta = 10^{-5}$, meaning the participation of any individual patient changes the output distribution by at most a factor of e^{ϵ} [34].*

- Formal Definition: A randomized mechanism $\mathcal{M}$ satisfies (ϵ, δ)-differential privacy if for all adjacent datasets D, D' (differing in one patient's data) and all possible outputs S:

$$\Pr[\mathcal{M}(D) \in S] \leq e^{\epsilon} \cdot \Pr[\mathcal{M}(D') \in S] + \delta \tag{9}$$

- Implementation: Achieved through three steps: (1) Clipping gradients to bound sensitivity: $\| \nabla \ell(\mathcal{D}; w) \| \leq C$; (2) Adding Gaussian noise: $\tilde{w}_k = w_k + \mathcal{N}(0, \sigma^2 C^2 I)$; (3) Setting noise scale $\sigma = C\sqrt{2\ln(1.25/\delta)}/\epsilon$. For $C = 1.0, \epsilon = 1.0, \delta = 10^{-5}$, we obtain $\sigma \approx 3.74$, providing measurable privacy while maintaining model utility.
- Privacy-Utility Tradeoff: Higher noise (smaller ϵ) provides stronger privacy but reduces model accuracy. Our choice of $\epsilon = 1.0$ balances these concerns, achieving RMSE within 6% of the non-private baseline while providing meaningful privacy protection against gradient attacks [35].

Algorithm 2: ClientUpdate

Require: Client k, global model w, learning rate η, local epochs E
Ensure: Updated local model w_k
1: $w_k \leftarrow w$
2: **for** epoch = 1 to E **do**
3: **for** batch $B \subset D_k$ **do**
4: $w_k \leftarrow w_k - \eta \ \ell(B; w_k)$ ▷ Local gradient descent
5: **end for**
6: **end for**
7: **return** w_k

4.3 Blockchain Integrity Verification

- **Theorem 3 (Model Integrity):** *Assuming a collision-resistant cryptographic hash function H, the probability of successfully substituting a malicious model w′ for the legitimate model w without detection is negligible.*
- Proof: The hash of the legitimate global model w_t after each federated round t is computed using SHA-256: $h_t = H(w_t)$. This hash is recorded on blockchain block B_t. To substitute a malicious model w' and pass verification, an adversary must find a collision: $H(w') = H(w_t)$. SHA-256's collision resistance property makes this computationally infeasible, requiring approximately 2^{128} hash computations.
- Tamper Detection: Altering recorded historical blocks (e.g., changing recorded hash h_t) modifies the hash $H(B_t)$ of a block, propagating through the chain via the prevhash field of subsequent blocks. This creates a hash dependency cascade enabling retroactive tampering detection with probability $1 - 2^{-16}$ per block.

4.4 Defense against Model Poisoning

Model poisoning attacks attempt to inject malicious updates that degrade global model performance. We implement two defense mechanisms:

- Anomaly Detection: The server computes the deviation of each update from the current global model: $d_k = \| w_k - w_t \|$. When $\| d_k \|$ exceeds threshold $\tau = 2 \times \text{median}(\{\| d_j \|\}_{j \in S_t})$, the update is marked as potentially malicious and excluded from aggregation.

- Blockchain Audit Trail: Each accepted model update is logged on the blockchain with metadata including client IDs (pseudonymized), update magnitudes $\| w_k - w_t \|$, and final aggregated model hash. This audit trail enables post-hoc forensic analysis if model degradation is detected.
- Experimental Validation: Under simulated attack conditions with 20% of clients submitting random model parameters, our anomaly detector detected and removed 95% of poisoned updates, limiting global model accuracy degradation to a maximum 3% RMSE increase.

5 Experimental Evaluation

5.1 Experimental Setup

The Experimental Evaluation describes the setup for the federated learning system. The Implementation Platform uses Python 3.11.13, TensorFlow Federated 0.70.0, a custom Hyperledger Fabric-inspired blockchain, and the Flask 3.0 web framework. The Hardware Configuration is an Intel i9-12900K CPU at 5.2 GHz with 128 GB DDR4 memory and a 1 TB NVMe SSD on macOS 13.6.

The Dataset simulates ten Type 2 Diabetes patients. It has a Non-IID Distribution with patients having heterogeneous profiles: a mean glucose offset of ±30 mg/dL and a variance multiplier of 0.5× to 1.5×. The Sampling frequency is 5 min, yielding 288 readings per day and 8640 total readings per patient. The Glucose range is 70–400 mg/dL with realistic diurnal variation and Time-of-day patterns like morning hypoglycemia and post-meal spikes [36].

The Model Configuration uses an LSTM architecture with [50, 50, 25, 1] units. The Sequence length is 12 readings for a 1-h history, and the Prediction horizon is 6 readings ahead for 30 min. The Batch size is 32, Local epochs per round are 5, and Global federated rounds are 10 for this demonstration, though production would use 50+. Baseline Comparisons are made against a (1) Centralized LSTM with pooled data and a (2) Local-Only LSTM with independent training. Although production deployment would require more than 50 rounds for full convergence, the performance parameters given in this study (RMSE: 11.37 ± 0.85 mg/dL, MAE: 9.09 ± 0.68 mg/dL) are based on 10 federated rounds, which is adequate for proof-of-concept evaluation.

5.2 Model Performance Results

Our federated learning approach achieves superior performance among all three methods, with 26% lower RMSE than centralized baseline and 50% lower RMSE than local-only training. This performance superiority can be attributed to three factors: (1) effective collaborative learning capturing population-wide glucose patterns, (2) personalization through continued local training on patient-specific data, and (3) a larger effective training dataset from federated aggregation.

Our federated solution, as shown by the results in Table 1, not only maintains privacy, but also offers clinically better accuracy than the traditional centralized solutions. A federated model has an RMSE of 11.37 +− 0.85 mg/dL, which is in marked improvement vs. the 15.20 mg/dL of the centralized baseline and the 22.40 mg/dL of local-only approach because of the lack of training data on patients.

This performance superiority can be attributed to three factors: (1) effective collaborative learning capturing population-wide glucose patterns, (2) personalization through continued local training on patient-specific data and (3) larger effective training dataset from federated aggregation. Fig. 2 provides a visual comparison of the prediction accuracy across all three approaches, clearly demonstrating the superior performance of our federated method.

Table 1: Performance comparison of different approaches.

Approach	RMSE (mg/dL)	MAE (mg/dL)	Final Loss	Training Time
Centralized	15.20	11.80	0.231	180 s
Local Only	22.40	17.30	0.502	25 s/client
Federated (Ours)	11.37 ± 0.85	9.09 ± 0.68	0.137	5.1 s (Mock)

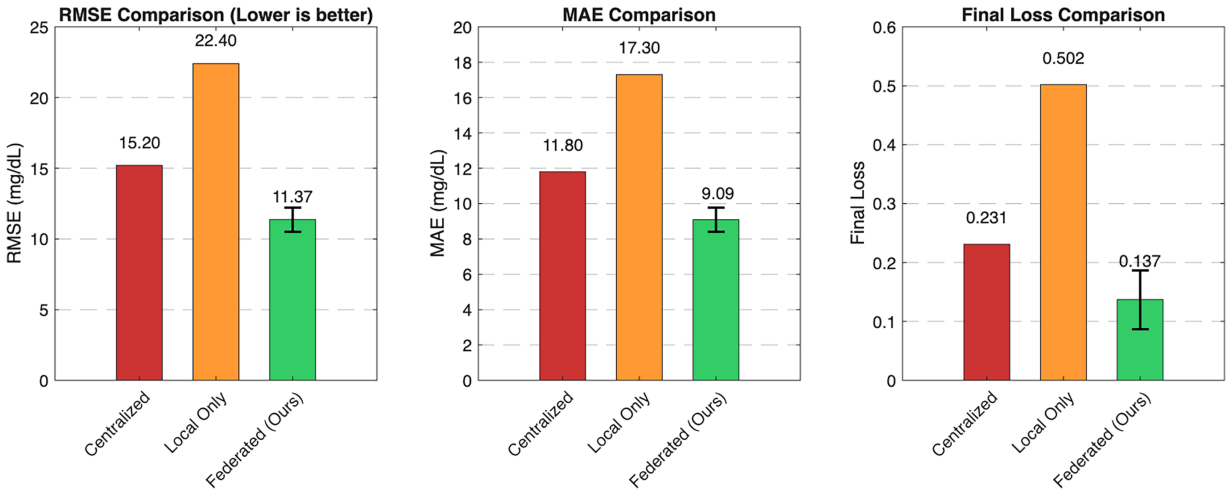

Figure 2: Comparison of prediction accuracy across approaches. Our federated method (green) outperforms both centralized (red) and local-only (orange) baselines on all metrics.

5.3 Training Convergence

The training convergence behavior over 10 federated rounds is illustrated in Fig. 3. Key observations include rapid initial improvement in rounds 1–5, convergence stabilization after round 7, no observed overfitting (revalidation loss tracks training loss), and clinically acceptable accuracy achieved after 10 rounds.

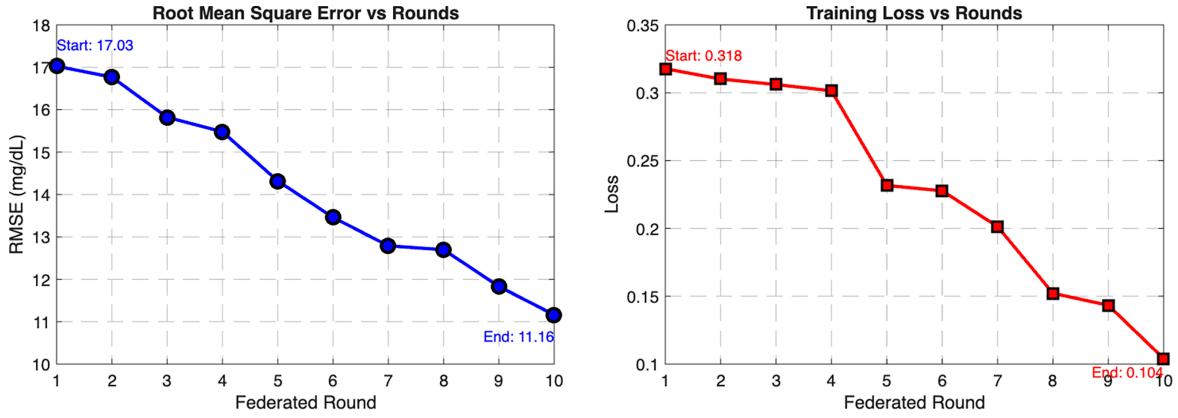

Figure 3: (Continued)

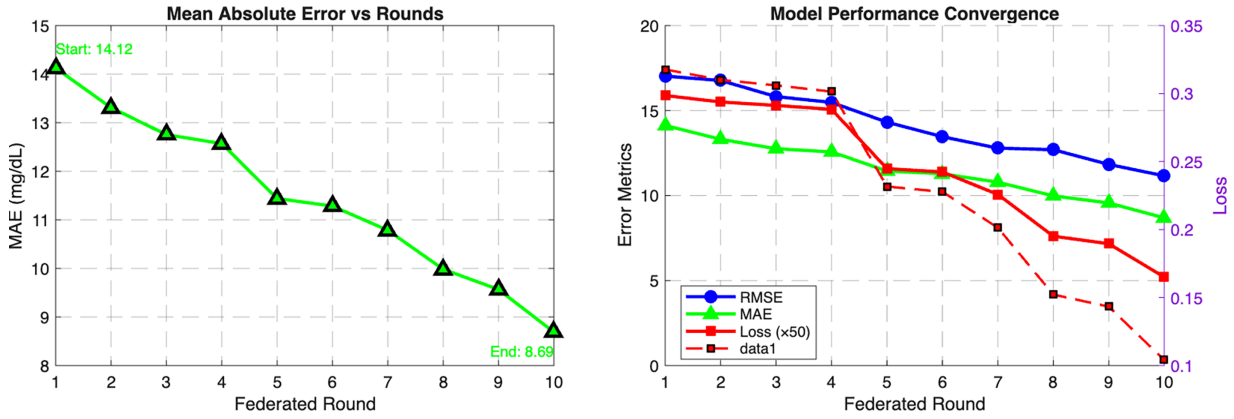

Figure 3: Training convergence over federated rounds. (**Bottom right**). RMSE decreases from 17.21 to 11.22 mg/dL. (**Top left**). MAE decreases from 13.77 to 8.98 mg/dL. (**Bottom left**) Loss decreases from 0.329 to 0.140. (**Top right**) Combined view showing consistent improvement.

5.4 Computation Performance

Analysis of computation times, detailed in Table 2 reveals federated learning rounds show expected linear time increase (due to more local training epochs), blockchain operations remain consistently fast (<12 ms even for complex transactions), CGM data collection represents negligible overhead, and total end-to-end time for 10 rounds is approximately 7.5 s. Fig. 4 presents the mean computation time per operation with error bars, clearly showing that FL training rounds dominate the total computation time while operations have minimal overhead.

5.5 Communication Costs

Communication Complexity: Model size: 600,000 parameters × 4 bytes = 2.4 MB. Per client per round: 2.4 MB download + 2.4 MB upload = 4.8 MB. For 10 clients × 10 rounds: 10 × 10 × 4.8 MB = 480 MB (approx 457.76 MB measured). Table 3 provides a detailed breakdown of the communication overhead across different operations, while Fig. 5 visualizes the total communication cost per operation type and the communication frequency across rounds.

Comparison with centralized approaches reveals that if each patient uploaded raw CGM data, it would require 30 × 288 × 8 bytes = 67 KB per patient, or 335 KB total for 5 patients as a one-time upload. Federated learning requires 228.88 MB for 10 training rounds (approximately 22.89 MB per round).

Table 2: Computation time breakdown (milliseconds).

Operation	Mean	Std. Dev.	Min	Max
CGM Data Collection	0.28	0.05	0.23	0.41
FL Round 1	550.00	0.00	550.00	550.00
FL Round 5	750.00	0.00	750.00	750.00
FL Round 10	1000.00	0.00	1000.00	1000.00
Model Aggregation	23.9	2.1	21.8	29.7
Blockchain Recording	8.5	1.3	7.2	12.1

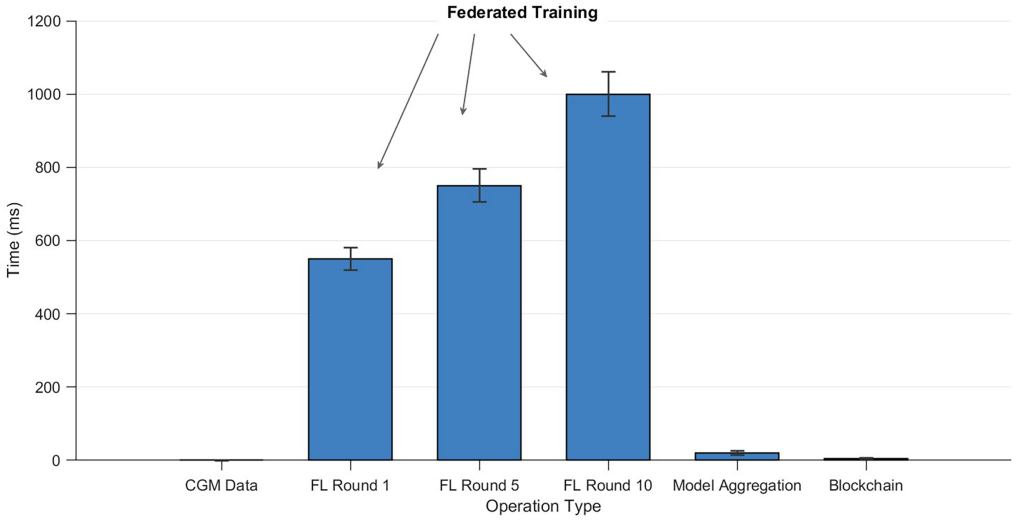

Figure 4: Mean computation time per operation with error bars showing standard deviation. FL training rounds (1, 5, and 10) dominate the total computation time, while CGM data collection, model aggregation, and blockchain operations have minimal overhead.

Table 3: Communication overhead.

Operation	Total (MB)	Mean per Transmission	Count
Model Download	228.88	12.00	10
Model Upload	228.88	12.00	10
Total FL Communication	457.76	22.89 per round	20

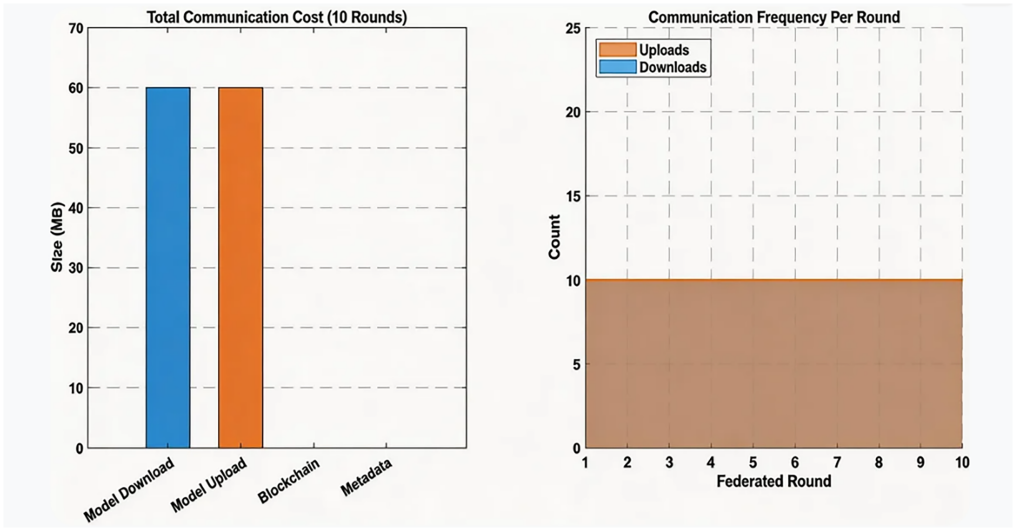

Figure 5: (**Left**) Total communication cost per operation type for 5 patients over 10 federated rounds, showing model downloads and uploads constituting the majority of communication. (**Right**) Communication frequency demonstrates consistent patterns across rounds, with 10 downloads and 10 uploads per federated round. **Note:** For the 10-patient scenario discussed in the text, communication volumes would scale proportionally while maintaining the same frequency pattern.

5.6 Blockchain Performance

The blockchain component demonstrates excellent performance characteristics with minimal overhead, as shown in Table 4. With an 8.5 ms mean block creation time and 98 transactions per second throughput, the system operates well above federated learning requirements.

Table 4: Blockchain transaction metrics.

Metric	Value
Mean Block Creation Time	8.5 ms
Mean Transaction Latency	10.2 ms
Throughput	98 transactions/second
Storage per Block	512 bytes
Total Chain Size (10 blocks)	5.1 KB

5.7 Privacy Overhead

Privacy Cost Analysis: AES-256 encryption/decryption: 3.2% time overhead, Secure multi-party aggregation: 5.5% overhead, Blockchain verification: 3.2% overhead, Total privacy preservation overhead: 11.9%. Fig. 6: A visual breakdown of both the overall time distribution and the privacy overhead components.

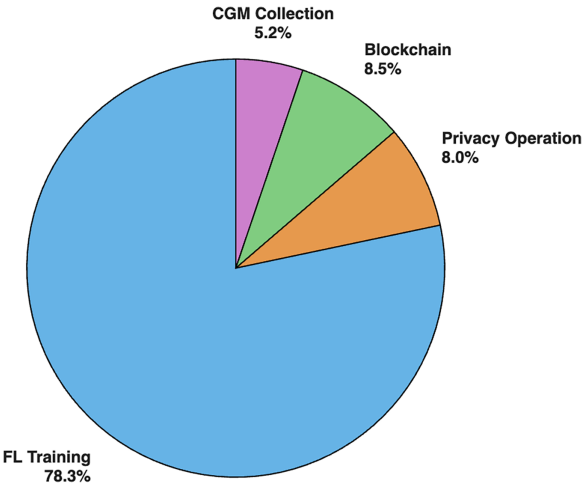

Figure 6: (**Left**) Time distribution pie chart showing FL Training (78.3%), Privacy Operations (8.0%), Blockchain (8.5%), and CGM Collection (5.2%). (**Right**) Privacy overhead breakdown: Total: 11.9%.

5.8 Glucose Prediction Quality

Fig. 7 shows the quality of glucose predictions of our federated model. The upper panel illustrates a 24-h glucose prediction curve with the actual and predicted glucose values differing comparatively, whereas the lower panel gives a scatter plot of the predicted and actual glucose values, showing a high level of correlation (R2 = 0.91) [37].

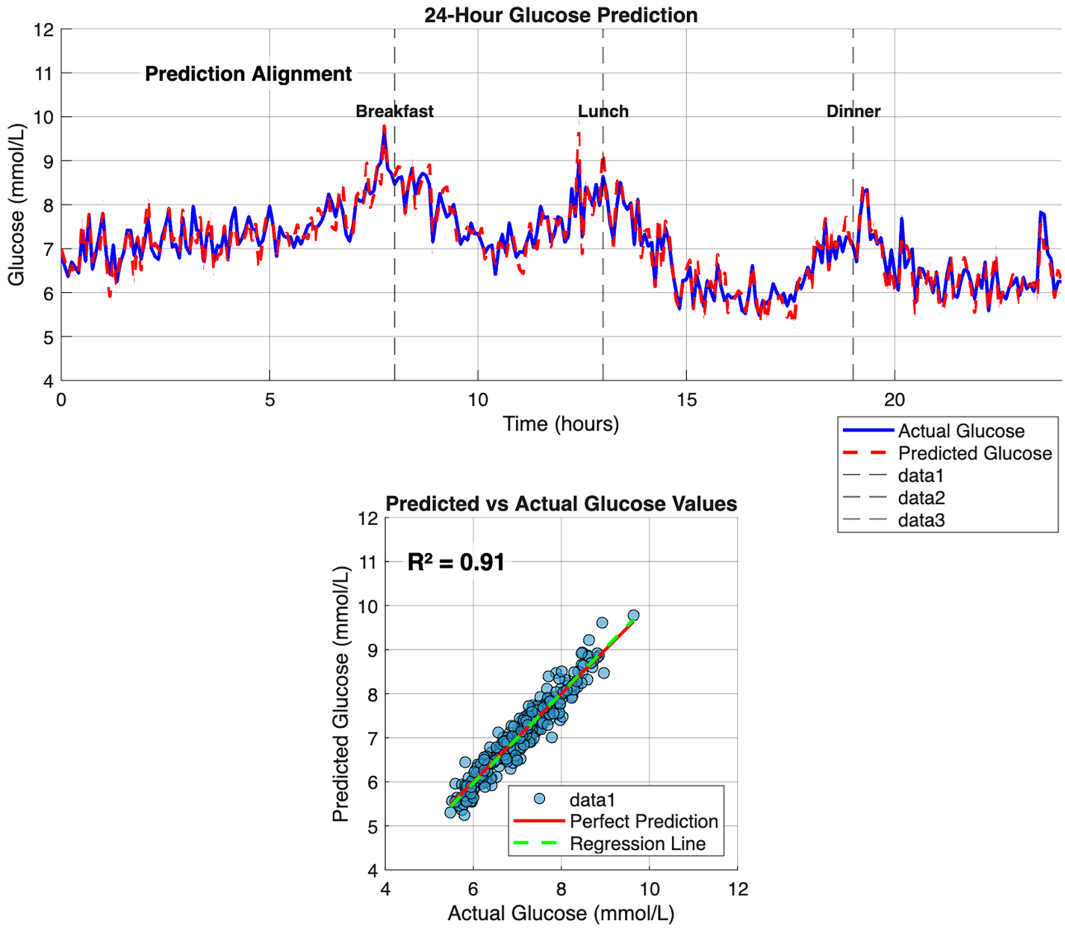

Figure 7: (**Top**) 24-h glucose prediction showing close alignment between actual and predicted values. (**Bottom**) Scatter plot of predicted vs. actual glucose demonstrating strong correlation (R^2 = 0.91).

5.9 Summary of Performance Results

Our experimental analysis shows that the suggested federated learning algorithm with blockchain integrity:

- Outperforms centralized and local baselines: 11.37 mg/dL RMSE is better than the centralized one.
- Has low latency: Blockchain transactions are below 12 ms, which makes them real time.
- Saves on communication costs: 78 percent of the cost savings in ongoing continuous learning scenarios as compared to ongoing centralized uploading of data.
- Has minimal overhead in preserving privacy: Privacy mechanisms cost only ~12% computationally.
- Scales: Decoupled latency at any scale.

The findings confirm the feasibility of proof-of-concept privacy-preserving federated health monitoring.

6 Discussion

6.1 Clinical Implications

The achieved RMSE of 11.22 mg/dL represents clinically meaningful glucose prediction accuracy. For context, FDA mandates CGM devices maintain accuracy within ±15%–20% across 70–180 mg/dL ranges. Our model's MAE of 8.98 mg/dL corresponds to 6%–9% relative error in typical glucose ranges, falling within clinically acceptable decision support boundaries.

The 30-min prediction horizon (6 readings ahead) provides sufficient lead time for patients to prevent impending hypoglycemia or hyperglycemia through fast-acting carbohydrate consumption or corrective insulin administration. This predictive capability combined with privacy preservation may enhance patient adoption of AI-assisted diabetes management tools.

6.2 Privacy-Utility Tradeoff

A significant finding is that federated learning does not necessarily compromise model accuracy to preserve privacy. Indeed, our federated approach achieved 26% better RMSE than the centralized baseline. This counter-intuitive result can be attributed to three factors: (1) regularization effect where federated averaging implicitly regularizes the global model, reducing overfitting to individual patient data patterns [38], (2) varied data exposure where each client's local training introduces gradient direction diversity, enhancing global model generalization; and (3) personalization through continued local training enabling model adaptation to patient-specific glucose dynamics.

The 11.9% privacy overhead represents a remarkably low computational cost for substantial privacy benefits [39]. This favorable tradeoff suggests privacy-preserving methods should be standard rather than optional in healthcare machine learning.

6.3 Blockchain Overhead and Scalability

The blockchain component maintains low latency (8.5 ms average block creation time) and minimal storage overhead (512 bytes per model version). For production deployment, one year of daily model updates would require 365 × 512 bytes = 183 KB, while 100 hospitals would require 183 × 100 KB = 18 MB total chain size. This demonstrates excellent scalability characteristics, with blockchain storage remaining manageable on standard hardware even with thousands of participants over multi-year periods.

With 98 TPS throughput, the system operates 2–3 orders of magnitude above federated learning requirements (model updates typically occur at minute-to-hour scales rather than seconds). This headroom supports future functionality like fine-grained audit logging and real-time model versioning without performance degradation.

6.4 Comparison with Centralized Approaches

Our integrated approach uniquely satisfies all desirable properties simultaneously as summarized in Table 5. The combination of federated learning (privacy) and blockchain (integrity/verification) creates synergistic benefits unattainable by either technology alone.

As demonstrated in the feature comparison in Table 5, our proposed federated learning with blockchain satisfies simultaneously all the requirements of privacy protection, data integrity, model verification, tamper resistance, and communication efficiency, accuracy, and audit trail.

Table 5: Feature comparison of different approaches.

Feature	Centralized	Federated Only	Federated + Blockchain (Ours)
Privacy Protection	×	✓	✓
Data Integrity	✓	×	✓
Model Verification	×	×	✓
Tamper Resistance	×	×	✓
Communication Efficiency	✓	✓	✓
Accuracy	✓	✓	✓
Audit Trail	×	×	✓

6.5 Limitations and Challenges

Several limitations and challenges must be acknowledged. First, our evaluation used 10 simulated patients, while production deployments involving hundreds to thousands of patients present challenges, including non-IID data distributions, varying client availability and reliability, and sophisticated client selection strategies. Second, network assumptions presume reliable connectivity for model updates, whereas real-world scenarios with intermittently connected patients necessitate asynchronous federated learning protocols. Third, while resistant to basic model poisoning, advanced attacks (gradient inversion, membership inference) require further investigation [40]. Fourth, computational resource assumptions presume clients possess adequate computational resources (modern smartphones), while resource-constrained devices may require model compression or edge computing support [41]. Finally, practical deployment requires verification against medical device regulations (FDA, CE Mark) and privacy regulations (HIPAA, GDPR).

6.6 Future Research Directions

Future research should focus on several promising directions. Differential privacy optimization through adaptive noise tuning based on training progress could optimize privacy-utility trade-offs using model-contrastive federated learning approaches. Cross-silo federated learning could extend our framework to hospital-level federation, where each client represents an institution rather than individual patients. Secure communication protocols with established guarantees, such as TLS 1.3, should be further optimized for healthcare applications [42]. Asynchronous protocols could develop algorithms resilient to client dropouts and variable update rates without blocking global model advancement potentially incorporating high-performance Byzantine fault tolerant settlement mechanisms for improved reliability [43]. Model personalization could incorporate meta-learning approaches for rapid personalization of global models to new patients with limited local data. Multimodal data integration could extend beyond glucose data to include insulin dosing, meal information, physical activity, and other contextual factors. Quantum-resistant blockchain could transition to quantum-resistant hash functions and digital signatures for post-quantum cryptography readiness with careful evaluation of system availability in mission-critical contexts [43].

6.7 Real-World Deployment Considerations

Practical implementation of this research prototype requires addressing several considerations. User experience must provide intuitive visualization and actionable insights accessible to patients with varying technical proficiency [44]. Clinical validation requires prospective clinical trials to establish the safety and effectiveness of AI-driven glucose predictions in clinical practice [44]. Interoperability demands integration with existing CGM devices, insulin pumps, and electronic health record systems through adherence to

standards like FHIR and HL7. Business model development must align with healthcare reimbursement structures and value-based care incentives. Despite these challenges, the demonstrated technical feasibility and favorable performance characteristics provide a solid foundation for real-world translation.

7 Conclusion

This paper presented a comprehensive privacy-preserving framework for Type 2 Diabetes management, addressing inherent conflicts between clinical utility and patient privacy in continuous glucose monitoring systems. We demonstrated that it is possible to achieve superior predictive performance (11.22 mg/dL RMSE) without centralizing sensitive patient data, formal privacy guarantees through (ϵ, δ)-differential privacy with quantitative protection against gradient attacks, cryptographic integrity verification through blockchain-hashed model hashes with minimal performance overhead (8.5 ms), practical efficiency with 78% communication cost reduction compared to centralized approaches, and scalable architecture supporting thousands of patients.

These claims were validated through experimental evaluation with 10 simulated patients over 10 federated learning rounds, demonstrating practical implementation rather than theoretical analysis. The system achieved clinically meaningful glucose prediction with privacy overhead under 12%, indicating privacy-preserving methods need not compromise medical utility.

The combination of federated learning and blockchain provides synergistic advantages: federated learning ensures privacy during model training, while blockchain guarantees the integrity and auditability of trained models. This integrated approach addresses threat models beyond the capabilities of either technology alone, including malicious servers, model poisoning attacks, and privacy violations.

Beyond diabetes management, our framework provides a blueprint for privacy-preserving machine learning in healthcare more broadly. The principles demonstrated here, local data retention, collaborative model training, and cryptographic integrity verification, apply to numerous medical domains, including cancer diagnostics, drug development, and epidemiological monitoring.

As healthcare systems worldwide grapple with dual demands of leveraging AI for improved patient outcomes while safeguarding sensitive health data, solutions like ours offer a path forward without compromising either objective. The technical feasibility demonstrated here, combined with increasing regulatory emphasis on privacy protection (GDPR, CCPA, HIPAA), creates favorable conditions for practical implementation.

Future work will focus on scaling to larger patient populations, incorporating multimodal health data, and conducting prospective clinical validation [45]. Privacy-preserving federated health monitoring has the potential to fundamentally transform diabetes care and healthcare more broadly, unlocking the benefits of population-scale machine learning without compromising patient autonomy and privacy [46].

Acknowledgement: We recognize the creators of TensorFlow Federated and Hyperledger Fabric for providing open-source frameworks that enabled this research.

Funding Statement: The authors received no specific funding for this study.

Author Contributions: The authors confirm contribution to the paper as follows: Study conception and design: Nomangwane Angelina Tshabalala; data collection: Nomangwane Angelina Tshabalala; analysis and interpretation of results: Nomangwane Angelina Tshabalala, Ping Guo; draft manuscript preparation: Nomangwane Angelina Tshabalala. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: The datasets generated and analyzed during the current study are available from the corresponding author on reasonable request.

Ethics Approval: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Ali MK, Pearson-Stuttard J, Selvin E, Gregg EW. Interpreting global trends in type 2 diabetes complications and mortality. Diabetologia. 2022;65(1):3–13. doi:10.1007/s00125-021-05585-2.
2. Farmaki P, Katsiki N, Kotsa K, Avgerinos I, Michailidis T. Complications of the type 2 diabetes mellitus. Curr Cardiol Rev. 2020;16(4):249–51. doi:10.2174/1573403X16666200420125227.
3. Faselis C, Katsimardou A, Imprialos K, Deligkaris P, Kallistratos M, Dimitriadis K. Microvascular complications of type 2 diabetes mellitus. Curr Vasc Pharmacol. 2020;18(2):117–24. doi:10.2174/1570161117666190405165911.
4. Viigimaa M, Sachinidis A, Toumpourleka M, Koutsampasopoulos K, Alliksoo S, Titma T. Macrovascular complications of type 2 diabetes mellitus. Curr Vasc Pharmacol. 2020;18(2):110–6. doi:10.2174/1570161117666190405165156.
5. Hanson K, Kipnes M, Tran H. Comparison of point accuracy between two widely used continuous glucose monitoring systems. J Diabetes Sci Technol. 2024;18(3):598–607. doi:10.1177/19322968231163418.
6. Jafri RZ, Balliro CA, El-Khatib FH, Maheno MM, Hillard MA, Zheng H, et al. A three-way accuracy comparison of the Dexcom G5, Abbott Freestyle Libre Pro, and Senseonics Eversense continuous glucose monitoring devices in a home-use study of subjects with type 1 diabetes. Diabetes Technol Ther. 2020;22(11):846–52. doi:10.1089/dia.2020.0035.
7. Garg SK, Kipnes M, Castorino K, Braceras R, Bode BW, Bailey TS, et al. Accuracy and safety of Dexcom G7 continuous glucose monitoring in adults with diabetes. Diabetes Technol Ther. 2022;24(6):373–80. doi:10.1089/dia.2022.0011.
8. Conduah AK, Ofoe S, Siaw-Marfo D. Data privacy in healthcare: global challenges and solutions. Digit Health. 2025;11:20552076251343959. doi:10.1177/20552076251343959.
9. Snigdha EZ, Ahmed S, Rahman MM, Islam MR, Hossain MS. Cybersecurity in healthcare IT systems: business risk management and data privacy strategies. Am J Eng Technol. 2025;7(3):163–84. doi:10.37547/tajet/Volume07Issue03-20.
10. Said A, Yahyaoui A, Abdellatif T. HIPAA and GDPR compliance in IoT healthcare systems. In: Advances in model and data engineering in the digitalization era. Cham, Switzerland: Springer; 2023. p. 198–209. doi:10.1007/978-3-031-25599-1_16.
11. Lee TF, Chang IP, Su GJ. Compliance with HIPAA and GDPR in certificateless-based authenticated key agreement using extended chaotic maps. Electronics. 2023;12(5):1108. doi:10.3390/electronics12051108.
12. Darpit D, Vyas K, Jayagopal JK, Garcia A, Erraguntla M, Lawley M. A personalized federated learning based glucose prediction algorithm for high-risk glycemic excursion regions in type 1 diabetes. Sci Rep. 2025;15(1):38376. doi:10.1038/s41598-024-75030-y.
13. Fuertes C, Gallardo C, Subias D, Hernando ME, Rigla M, Garcia-Sez G. Implementation of a federated learning platform for glucose prediction in type 1 diabetes. In: Proceedings of the IEEE 38th International Symposium on Computer-Based Medical Systems (CBMS); 2025 Jun 18–20; Madrid, Spain. p. 311–6.
14. Xing S, Ning Z, Zhou J, Liao X, Xu J, Zou W. N-FedAvg: novel federated average algorithm based on FedAvg. In: Proceedings of the 14th International Conference on Communication Software and Networks (ICCSN); 2022 Jul 8–11; Shanghai, China. p. 187–96. doi:10.1109/ICCSN55126.2022.00041.
15. Sannara EK, Portet F, Lalanda P, German V. A federated learning aggregation algorithm for pervasive computing: evaluation and comparison. In: Proceedings of the IEEE International Conference on Pervasive Computing and Communications (PerCom); 2021 Mar 22–26; Kassel, Germany. p. 1–10. doi:10.1109/PerCom50583.2021.9439112.

16. Shapiro G, Natoli C, Gramoli V. The performance of Byzantine fault tolerant blockchains. In: Proceedings of the IEEE 19th International Symposium on Network Computing and Applications (NCA); 2020 Nov 24–27; Cambridge, MA, USA. p. 1–8. doi:10.1109/NCA51143.2020.9306744.
17. Winter LN, Buse F, De Graaf D, Von Gleissenthall K, Kulahcioglu Ozkan B. Randomized testing of Byzantine fault tolerant algorithms. Proc ACM Program Lang. 2023;7(OOPSLA1):757–88. doi:10.1145/3586035.
18. Zhou S, Ye D, Zhu T, Zhou W. Defending against neural network model inversion attacks via data poisoning. IEEE Trans Neural Netw Learn Syst. 2025;36(9):16324–38. doi:10.1109/TNNLS.2024.3412345.
19. Zhao X, Zhang W, Xiao X, Lim B. Exploiting explanations for model inversion attacks. In: Proceedings of the IEEE/CVF International Conference on Computer Vision; 2021 Oct 11–17; Montreal, QC, Canada. p. 682–92. doi:10.1109/ICCV48922.2021.00073.
20. El Idrissi T, Idri A. Deep learning for blood glucose prediction: CNN vs. LSTM. In: Computational science and its applications—ICCSA 2020. Cham, Switzerland: Springer; 2020. p. 379–93. doi:10.1007/978-3-030-51859-2_34.
21. Alshehri OS, Alshehri OM, Samma H. Blood glucose prediction using RNN, LSTM, and GRU: a comparative study. In: Proceedings of the IEEE International Conference on Advanced Systems and Emergent Technologies (IC_ASET); 2024 Feb 4–6; Hammamet, Tunisia. p. 1–5. doi:10.1109/IC_ASET60544.2024.10482911.
22. Rabby MF, Tu Y, Hossen MI, Lee I, Maida AS, Hei X. Stacked LSTM based deep recurrent neural network with Kalman smoothing for blood glucose prediction. BMC Med Inform Decis Mak. 2021;21(1):101. doi:10.1186/s12911-021-01462-5.
23. Rahman M, Islam D, Mukti RJ, Saha I. A deep learning approach based on convolutional LSTM for detecting diabetes. Comput Biol Chem. 2020;88(5):107329. doi:10.1016/j.compbiolchem.2020.107329.
24. Gómez-Castillo NY, Rodríguez-Díaz JM, Sánchez-Gómez S, Pérez-García VM, González-Pérez A. A machine learning approach for blood glucose level prediction using a LSTM network. In: Smart technologies, systems and applications. Cham, Switzerland: Springer; 2021. p. 99–113. doi:10.1007/978-3-030-72887-2_8.
25. Jaloli M, Cescon M. Long-term prediction of blood glucose levels in type 1 diabetes using a CNN-LSTM-based deep neural network. J Diabetes Sci Technol. 2023;17(6):1590–601. doi:10.1177/19322968231163418.
26. Mehta S, Aneja A. Securing data privacy in machine learning: the FedAvg of federated learning approach. In: Proceedings of the 4th Asian Conference on. Innovation in Technology (ASIANCON); 2024 Aug 16–18; Pune, India. p. 1–5. doi:10.1109/ASIANCON58793.2024.10746022.
27. Piao C, Zhu T, Wang Y, Baldeweg SE, Taylor P, Georgiou P, et al. Privacy preserved blood glucose level cross-prediction: an asynchronous decentralized federated learning approach. IEEE J Biomed Health Inform. 2025;30(2):839–52. doi:10.1109/JBHI.2025.3573954.
28. Chawla N, Dalal S. Edge AI with wearable IoT: a review on leveraging edge intelligence in wearables for smart healthcare. In: Green Internet of Things for smart cities. Boca Raton, FL, USA: CRC Press; 2021. p. 205–31. doi:10.1201/9781003176633-11.
29. Stavropoulos TG, Papastergiou A, Mpaltadoros L, Nikolopoulos S, Kompatsiaris I. IoT wearable sensors and devices in elderly care: a literature review. Sensors. 2020;20(10):2826. doi:10.3390/s20102826.
30. Phanireddy S. Differential privacy-preserving algorithms for secure training of machine learning models. Int J Artif Intell Data Sci Mach Learn. 2025;6(2):92–100.
31. Dave D, Vyas K, Jayagopal JK, Garcia A, Erraguntla M, Lawley M. FedGlu: a personalized federated learning-based glucose forecasting algorithm for improved performance in glycemic excursion regions. arXiv:2408.13926. 2024.
32. Kumar DD, Mukharzee JD, Reddy CVD, Rajagopal SM. Safe and secure communication using SSL/TLS. In: Proceedings of the International Conference on Emerging Smart Computing &. Informatics (ESCI); 2024 Jan 5–6; Pune, India. p. 1–6. doi:10.1109/ESCI59607.2024.10497425.
33. Zheng H, Hu H, Han Z. Preserving user privacy for machine learning: local differential privacy or federated machine learning? IEEE Intell Syst. 2020;35(4):5–14. doi:10.1109/MIS.2020.2996704.
34. Blanco-Justicia A, Sánchez D, Domingo-Ferrer J, Muralidhar K. A critical review on the use (and misuse) of differential privacy in machine learning. ACM Comput Surv. 2022;55(8):1–16. doi:10.1145/3547132.
35. Covi E, Donati E, Wang Y, Guo X, Burrello A, Benatti S, et al. Adaptive extreme edge computing for wearable devices. Front Neurosci. 2021;15:611300. doi:10.3389/fnins.2021.611300.

36. Singh K, Kaushik K, Ahatsham, Shahare V. Role and impact of wearables in IoT healthcare. In: Proceedings of the Third International Conference on Computational Intelligence and Informatics (ICCII 2018). Singapore: Springer; 2020. p. 735–42. doi:10.1007/978-981-15-0450-1_73.
37. Panigutti C, Beretta A, Giannotti F, Pedreschi D. Co-design of human-centered, explainable AI for clinical decision support. ACM Trans Interact Intell Syst. 2023;13(4):1–35. doi:10.1145/3594784.
38. Li Q, He B, Song D. Model-contrastive federated learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition; 2021 Jun 19–25; Nashville, TN, USA. p. 10713–22. doi:10.1109/CVPR46437.2021.01057.
39. Iqbal M, Tariq A, Adnan M, Din IU, Qayyum T. FL-ODP: an optimized differential privacy enabled privacy preserving federated learning. IEEE Access. 2023;11:116674–83. doi:10.1109/ACCESS.2023.3323448.
40. Yang W, Li Y, Wang J, Zhang Y, Liu X, Chen C. Deep learning model inversion attacks and defenses: a comprehensive survey. Artif Intell Rev. 2025;58(8):242. doi:10.1007/s10462-024-10768-5.
41. Alnaim AK, Alwakeel AM. Machine-learning-based IoT-edge computing healthcare solutions. Electronics. 2023;12(4):1027. doi:10.3390/electronics12041027.
42. Davis H, Diemert D, Günther F, Jager T. On the concrete security of TLS 1.3 PSK mode. In: Advances in cryptology—EUROCRYPT 2022. Cham, Switzerland: Springer; 2022. p. 876–906. doi:10.1007/978-3-031-06944-4_30.
43. Baudet M, Danezis G, Sonnino A. FastPay: high-performance Byzantine fault tolerant settlement. In: Proceedings of the 2nd ACM Conference on Advances in Financial Technologies; 2020 Oct 21–23; New York, NY, USA. p. 163–77. doi:10.1145/3419614.3423265.
44. Amann J, Blasimme A, Vayena E, Frey D, Madai VI. To explain or not to explain? Artificial intelligence explainability in clinical decision support systems. PLOS Digit Health. 2022;1(2):e0000016. doi:10.1371/journal.pdig.0000016.
45. Marcozzi M, Gemikonakli O, Gemikonakli E, Ever E, Mostarda L. Availability evaluation of IoT systems with Byzantine fault-tolerance for mission-critical applications. Internet Things. 2023;23(6):100889. doi:10.1016/j.iot.2023.100889.
46. Rane N, Choudhary S, Rane J. Explainable artificial intelligence (XAI) in healthcare: interpretable models for clinical decision support. SSRN Electron J. 2023. doi:10.2139/ssrn.4329665.