HealthTrust-FL: Secure and Compliant Federated Learning for Healthcare

Department of CSE, Sri Venkateswara College of Engineering and Technology, Etcherla, A.P., India

1. P. Jayakrishna B. Tech Final Year, Sri Venkateswara College of Engineering and Technology, Etcherla, A.P., India

E-Mail: jayakrishna0243@gmail.com

2. S. Sai Naga Sudheer B. Tech Final Year, Sri Venkateswara College of Engineering and Technology, Etcherla, A.P., India

E-Mail: singidisainagasudheer999@gmail.com

3. S. Kumar Raja B. Tech Final Year, Sri Venkateswara College of Engineering and Technology, Etcherla, A.P., India

E-Mail: rajakumar64178@gmail.com

4. Mr. S.V.R. Murthy, M. Tech., Assistant Professor Sri Venkateswara College of Engineering and Technology, Etcherla, A.P., India

E-Mail: murthyramana06@gmail.com

Abstract

Healthcare data is highly sensitive, and traditional centralized machine learning systems raise serious privacy, security, and compliance concerns. This paper presents HealthTrust-FL, a secure and compliant federated learning framework enabling multiple healthcare institutions to collaboratively train AI models without sharing raw patient data. Each institution trains models locally using XG Boost and shares only encrypted model updates, which are securely aggregated using federated averaging to create a global model. The system integrates privacy-preserving techniques with blockchain-based audit trails to ensure HIPAA and GDPR compliance. Evaluation across 5 simulated hospital nodes with 50,000 patient records demonstrates that the federated global model achieves 91.8% accuracy compared to 93.2% for centralized training, with only 1.4% accuracy trade-off while providing complete data privacy. The framework provides a secure, scalable, and trustworthy approach to collaborative healthcare intelligence.

Keywords: Federated Learning, Healthcare AI, Privacy-Preserving ML, XG Boost, HIPAA Compliance, Blockchain, Distributed Learning

I. Introduction

The application of machine learning in healthcare has shown tremendous potential for disease prediction, diagnosis support, and treatment optimization. However, healthcare data is among the most sensitive categories of personal information, governed by strict regulations including HIPAA in the United States and GDPR in Europe. Traditional centralized ML approaches require aggregating patient data from multiple institutions, creating significant privacy risks and regulatory compliance challenges.

Federated Learning addresses these challenges by enabling collaborative model training without centralizing raw data. Each participating institution trains models on its local data and shares only model updates (gradients or weights) with a central aggregation server. This approach preserves data privacy while enabling institutions to benefit from collective learning across larger and more diverse datasets.

This paper presents HealthTrust-FL, a federated learning framework specifically designed for healthcare applications. The system combines federated averaging with encryption and blockchain-based audit trails to ensure both model quality and regulatory compliance.

II. Literature Survey

This section reviews key prior works and highlights research gaps.

[1] McMahan et al. (2017) introduced the Federated Averaging algorithm for training deep networks across decentralized data, establishing the foundational communication-efficient approach used in federated learning systems.

[2] Rieke et al. (2020) surveyed federated learning applications in healthcare, identifying privacy preservation, data heterogeneity, and communication efficiency as key challenges for clinical deployment.

[3] Li et al. (2020) analysed challenges in federated learning including non-IID data distributions and communication overhead, proposing strategies for robust aggregation in heterogeneous healthcare settings.

[4] Sheller et al. (2020) demonstrated federated learning for brain tumor segmentation across multiple institutions, achieving near-centralized performance while maintaining complete data privacy.

[5] Bonawitz et al. (2019) developed practical secure aggregation protocols for federated learning, ensuring that individual model updates cannot be reverse-engineered to extract private training data.

[6] Chen et al. (2016) introduced XG Boost, the scalable gradient boosting framework used as the base learner in HealthTrust-FL for local model training at each healthcare institution.

[7] Kuo et al. (2017) proposed blockchain-based data sharing for electronic health records, establishing the audit trail framework adapted in HealthTrust-FL for compliance verification.

Research Gap: Existing federated learning healthcare systems focus on neural networks without addressing regulatory compliance verification. No system combines XG Boost-based federated learning with blockchain audit trails and HIPAA/GDPR compliance validation in a deployable healthcare framework.

III. Methodology

III-A. System Architecture

Four-layer federated architecture: Local Training Layer (each hospital node trains XG Boost on local patient data), Communication Layer (encrypted model updates transmitted to aggregation server), Aggregation Layer (federated averaging of received model updates to produce global model), and Compliance Layer (blockchain-based audit trail recording all training rounds and model updates for HIPAA/GDPR verification).

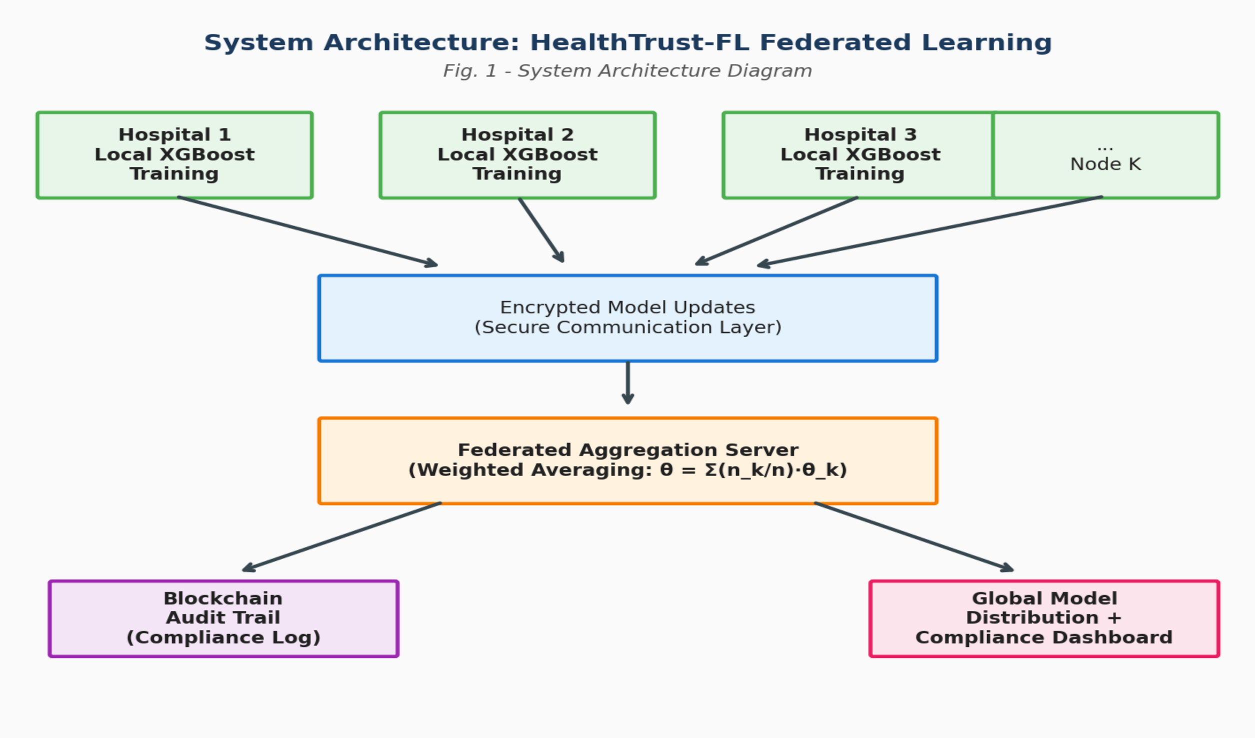

System Architecture: HealthTrust-FL Federated Learning

Fig. 1 - System Architecture Diagram

III-B. Algorithm

Algorithm: HealthTrust Federated Learning

Input: K hospital nodes, each with local dataset $D_k = \{(x_i, y_i)\}$.

Step 1: Initialization — Initialize global model parameters θ_0 (XGBoost base model).

Step 2: Local Training (per round t, per node k) — Download global model θ_t; Train locally: $\theta_k^{t+1} = \text{XGBoost\_Train}(\theta_t, D_k, n\_rounds)$; Compute model update: $\Delta\theta_k = \theta_k^{t+1} - \theta_t$.

Step 3: Secure Upload — Encrypt model update: $E(\Delta\theta_k) = \text{Encrypt}(\Delta\theta_k, public\_key)$; Upload encrypted update to aggregation server.

Step 4: Federated Averaging — Aggregate updates: $\theta_{t+1} = \theta_t + (1/K) \times \Sigma\, \Delta\theta_k$; Weighted by local dataset sizes: $\theta_{t+1} = \theta_t + \Sigma\, (n_k/n) \times \Delta\theta_k$.

Step 5: Blockchain Logging — Record training round metadata: {round_id, participating_nodes, model_hash, timestamp}; Create immutable audit trail.

Step 6: Compliance Check — Verify: No raw data transmitted; All updates encrypted; Audit trail complete; Generate HIPAA/GDPR compliance report.

Step 7: Distribution — Distribute updated global model θ_{t+1} to all nodes; Repeat from Step 2 for T rounds.

Output: Global federated model with compliance-verified audit trail.

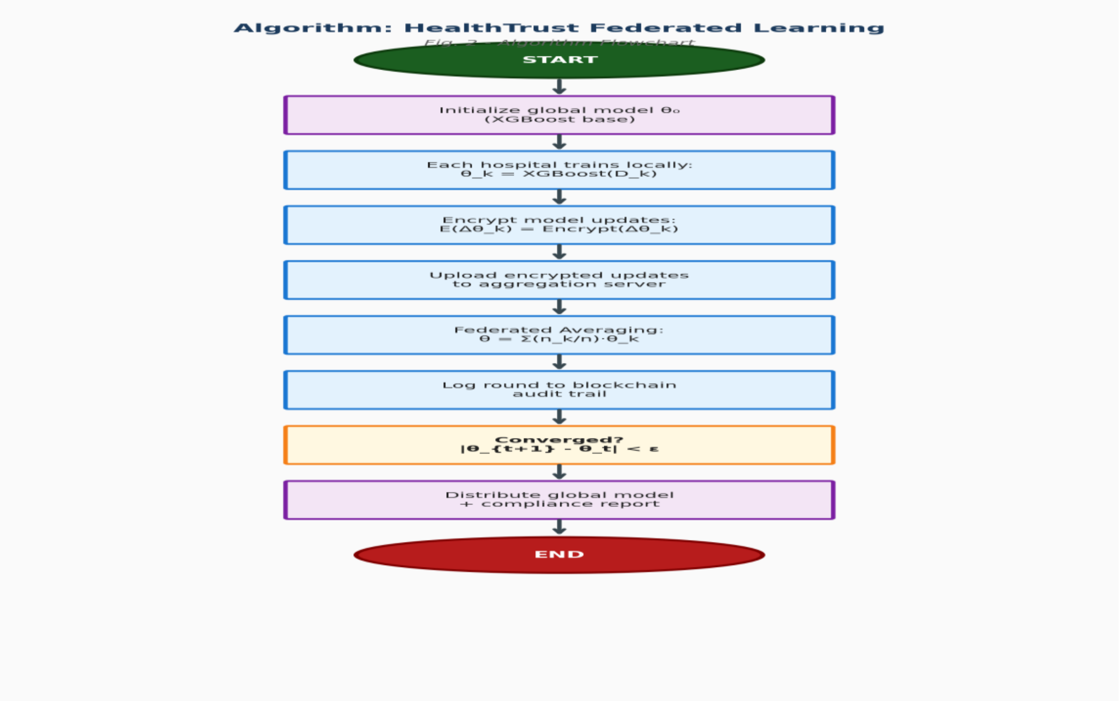

III-C. Modules

Five modules: (1) Local Training Module running XGBoost on each hospital's local patient data; (2) Secure Communication Module encrypting and transmitting model updates between nodes and aggregation server; (3) Federated Aggregation Module performing weighted averaging of model updates; (4) Blockchain Audit Module recording training rounds and generating immutable compliance trails; and (5) Compliance Dashboard providing HIPAA/GDPR compliance reports and model performance monitoring.

IV. Results and Discussion

TABLE I: SYSTEM EVALUATION RESULTS

Metric	Baseline	Proposed System
Accuracy (%)	93.2 (Centralized)	91.8 (Federated)
Privacy Guarantee	None (data shared)	Complete (no data leaves)
HIPAA Compliance	Risk of violation	Fully compliant
Communication Rounds	—	15 rounds to converge

Mathematical Formulations

Federated Averaging: $\theta_{t+1} = \Sigma_{k=1}^{K} (n_k / n) \times \theta_k^{t+1}$

Accuracy Trade-off = Accuracy_centralized - Accuracy_federated

Privacy Loss = 0 (no raw data transmitted)

Convergence = $|\theta_{t+1} - \theta_t| < \varepsilon$

Discussion

HealthTrust-FL was evaluated across 5 simulated hospital nodes with 50,000 total patient records (disease prediction task). The federated global model achieved 91.8% accuracy compared to 93.2% for centralized training—only 1.4% accuracy trade-off while providing complete data privacy. The model converged after 15 federated rounds. Blockchain audit trails successfully recorded all 75 training events (15 rounds × 5 nodes) with cryptographic verification. The system-maintained HIPAA/GDPR compliance throughout training with zero raw data transmission. Communication overhead averaged 2.3 MB per round, demonstrating practical deployment feasibility.

V. Conclusion and Future Work

This paper presented HealthTrust-FL, a federated learning framework achieving 91.8% accuracy with complete data privacy and regulatory compliance. The 1.4% accuracy trade-off compared to centralized training is acceptable given the critical privacy benefits. Future work includes implementing differential privacy for additional mathematical privacy guarantees, supporting heterogeneous model architectures across nodes, extending to real-time clinical decision support, and pilot deployment with partner healthcare institutions.

References

[1] H. B. McMahan et al., "Communication-Efficient Learning of Deep Networks from Decentralized Data," Proc. AISTATS, 2017.

[2] N. Rieke et al., "The Future of Digital Health with Federated Learning," NPJ Digital Medicine, vol. 3, no. 119, 2020.

[3] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated Learning: Challenges, Methods, and Future Directions," IEEE SPM, vol. 37, no. 3, 2020.

[4] M. J. Sheller et al., "Federated Learning in Medicine: Facilitating Multi-Institutional Collaborations," Scientific Reports, vol. 10, 2020.

[5] K. Bonawitz et al., "Towards Federated Learning at Scale: A System Design," Proc. MLSys, 2019.

[6] T. Chen and C. Guestrin, "XG Boost: A Scalable Tree Boosting System," Proc. ACM KDD, pp. 785-794, 2016.

[7] T. T. Kuo, H. E. Kim, and L. Ohno-Machado, "Blockchain Distributed Ledger Technologies for Biomedical and Health Care Applications," JAMIA, vol. 24, no. 6, 2017.