Enhanced Privacy-Preserving Blockchain-Enabled Federated Learning (EPP-BCFL) For Secure Healthcare Intelligence

S. Senthamarai[1], Dr. R. Mala[2]
[1]*Research Scholar, Ph.D. (Computer Science), Alagappa University*
[2]*Assistant Professor, Department of Computer Science, Alagappa University College of Arts and Science for Women Satellite Campus, Thondi*

***Abstract*—The rapid integration of Artificial Intelligence (AI), the Internet of Medical Things (IoMT), and Electronic Health Records (EHRs) has substantially transformed modern healthcare by enabling intelligent disease diagnosis, remote patient monitoring, and personalized clinical treatment. However, conventional centralized machine learning approaches require healthcare institutions to transmit sensitive patient data to centralized servers, thereby raising serious privacy, security, and regulatory concerns. Federated Learning (FL) has emerged as a promising distributed learning paradigm that facilitates collaborative model training without direct exchange of raw patient data. Despite these advantages, existing FL frameworks remain susceptible to model poisoning attacks, excessive communication overhead, insufficient trust management mechanisms, and malicious participant behavior.**

To address these persistent limitations, this paper proposes an Enhanced Privacy-Preserving Blockchain-Enabled Federated Learning (EPP-BCFL) framework designed for secure healthcare intelligence. The proposed architecture seamlessly integrates federated learning with blockchain technology through a lightweight Proof-of-Stake Byzantine Fault Tolerant (PoS-BFT) consensus mechanism, achieving improved trust, scalability, communication efficiency, and robustness against adversarial attacks. Smart contracts are employed to validate encrypted model updates, while secure aggregation protocols safeguard sensitive healthcare information throughout the collaborative training process. The proposed EPP-BCFL framework was rigorously evaluated using three real-world healthcare datasets: the MIMIC-III clinical database, the Breast Cancer Wisconsin Diagnostic Dataset, and the ChestX-ray14 medical imaging dataset. Experimental results demonstrate that the proposed framework achieves 95.2% classification accuracy, reduced communication overhead, enhanced privacy preservation, low blockchain validation latency, and strong resilience against adversarial attacks when compared with existing federated learning approaches.

***Index Terms*—Federated Learning, Blockchain, Healthcare Intelligence, Privacy Preservation, Electronic Health Records, IoMT, Smart Contracts, PoS-BFT, Deep Learning, Secure Healthcare Analytics.**

I. INTRODUCTION

The growing adoption of digital healthcare technologies — including Electronic Health Records (EHRs), cloud computing platforms, wearable devices, and IoMT sensors — has generated unprecedented volumes of healthcare data. Artificial Intelligence (AI) and deep learning models have demonstrated remarkable capabilities in disease diagnosis, medical image analysis, remote patient monitoring, and intelligent clinical decision support.

Traditional centralized machine learning systems require healthcare institutions to consolidate sensitive patient data on shared servers before model training can begin. Although this centralized paradigm achieves competitive predictive performance, it introduces severe privacy risks, including unauthorized data access, information leakage, targeted cyberattacks, and potential violations of regulatory frameworks such as HIPAA and GDPR.

Federated Learning (FL) offers an alternative by allowing multiple healthcare organizations to collaboratively train a global AI model while retaining patient data within their local infrastructure. Despite its inherent privacy advantages, conventional FL systems remain vulnerable to malicious model

updates, adversarial attacks, communication inefficiencies, and the absence of a robust decentralized trust management layer.

Blockchain technology, with its immutable distributed ledgers, decentralized verification mechanisms, smart contracts, and secure consensus protocols, provides a compelling foundation for building trust and transparency in collaborative learning environments. Integrating blockchain with federated learning enables cryptographically verifiable validation of model updates and significantly reduces the risk of model corruption by malicious participants.

The primary contributions of this work are summarized as follows:

- Development of a blockchain-enabled federated learning framework tailored for secure healthcare intelligence.
- Integration of a lightweight PoS-BFT consensus mechanism to achieve low-latency blockchain validation.
- Implementation of encrypted parameter exchange and secure aggregation to protect model updates during transmission.
- Real-time detection of malicious model updates through smart contract-based verification.
- Comprehensive evaluation using real-world healthcare datasets and systematic comparison against existing baseline approaches.

II. LITERATURE REVIEW

Recent progress in AI, distributed computing, blockchain technology, and healthcare informatics has accelerated the development of privacy-preserving collaborative learning systems. Federated Learning has received considerable attention as a decentralized machine learning paradigm enabling multi-institutional model training without direct data sharing. Researchers have explored diverse aspects of federated learning, blockchain integration, and secure healthcare intelligence to overcome the limitations inherent in centralized healthcare analytics.

McMahan et al. [1] introduced the Federated Averaging (FedAvg) algorithm, which became the cornerstone of modern federated learning. Their work showed that distributed clients could collaboratively train deep learning models while minimizing communication costs and preserving local data privacy. Building on this foundation, Kairouz et al. [2] conducted a comprehensive survey addressing the major challenges in federated learning, including communication overhead, data heterogeneity, scalability, and adversarial vulnerabilities.

Li et al. [3] examined the technical obstacles and future directions of federated learning systems, with particular focus on non-IID data distributions and optimization instabilities in decentralized environments. Yang et al. [4] complemented this analysis by discussing practical federated machine learning applications across healthcare, finance, and industrial domains, underscoring its relevance for privacy-sensitive use cases.

Privacy preservation remains a central challenge in collaborative healthcare learning. Bonawitz et al. [5] proposed a secure aggregation protocol that prevents the central server from accessing individual client updates, significantly improving the confidentiality of exchanged model parameters. Blockchain technology has gained growing attention for decentralized trust management. Nakamoto [6] introduced the foundational concept of immutable distributed ledgers through Bitcoin, while Zheng et al. [7] provided a comprehensive overview of blockchain architectures and consensus mechanisms. Xu et al. [8] further explored blockchain application architectures and their integration with distributed systems.

Androulaki et al. [9] introduced Hyperledger Fabric, a permissioned blockchain platform suitable for enterprise and healthcare environments, supporting smart contracts and distributed consensus. Singh and Chatterjee [10] proposed blockchain-based solutions to improve medical data protection, while Agbo et al. [11] conducted a systematic review highlighting blockchain's potential for secure medical record management and decentralized trust.

Healthcare-specific federated learning has also attracted growing research interest. Hussein et al. [12] investigated FL techniques for healthcare informatics, demonstrating their effectiveness in protecting sensitive medical data. Rieke et al. [13] emphasized FL's role in enabling multi-institutional AI collaboration without violating patient privacy regulations. Sheller et al. [14] validated FL for

distributed medical imaging, demonstrating competitive accuracy while preserving confidentiality. The challenge of data heterogeneity in federated learning was examined by Li et al. [15], who studied the convergence behavior of FedAvg under Non-IID distributions and highlighted its impact on model performance. Regarding consensus mechanisms, King and Nadal [16] introduced Proof-of-Stake (PoS) as an energy-efficient alternative to Proof-of-Work, while Castro and Liskov [17] proposed Byzantine Fault Tolerant (BFT) consensus for improved fault tolerance. Combining PoS and BFT offers enhanced scalability, reduced latency, and resilience against malicious participants.

Several recent studies have explored blockchain-federated learning integration for healthcare security. Zhou et al. [18] proposed a blockchain-based FL framework demonstrating improved trust management, while Kumar et al. [19] developed a distributed deep learning architecture enhancing privacy and secure collaborative analytics. Alazab et al. [20] further demonstrated improvements in privacy, communication efficiency, and attack resistance. However, existing approaches continue to face challenges related to scalability, blockchain latency, communication overhead, and computational complexity.

Motivated by these limitations, the proposed EPP-BCFL framework integrates federated learning with a lightweight PoS-BFT blockchain mechanism to simultaneously improve scalability, communication efficiency, model integrity, and adversarial attack resistance, while maintaining high classification accuracy and strong privacy preservation in intelligent healthcare systems.

III. PROPOSED EPP-BCFL METHODOLOGY

The proposed EPP-BCFL framework integrates federated learning with blockchain technology to establish a secure, decentralized, and privacy-preserving healthcare intelligence ecosystem. The architecture enables multiple healthcare institutions — including hospitals, diagnostic laboratories, clinics, and IoMT-enabled healthcare centers — to collaboratively train AI models without sharing sensitive patient information. The framework is organized into three major layers: the Healthcare Edge Layer, the Federated Learning Layer, and the Blockchain Trust Layer. Together, these layers ensure secure data processing, trusted model aggregation, and decentralized verification of healthcare intelligence. Figure 1 illustrates the overall system architecture.

3.1. Data Collection and Local Management

Healthcare data are collected from distributed organizations and IoMT devices, encompassing EHRs, medical imaging data, laboratory reports, physiological sensor measurements, and patient demographic information. Critically, the proposed framework ensures that all patient data remain locally stored within each institution, thereby preserving privacy and ensuring compliance with HIPAA and GDPR. The healthcare edge nodes independently manage their local datasets and participate in collaborative federated learning without exposing raw medical records.

3.2. Data Preprocessing

Prior to model training, each healthcare institution applies local preprocessing to enhance data quality and training efficiency. This stage encompasses missing value imputation, Min-Max normalization, noise removal, feature extraction, image resizing, data anonymization, and augmentation techniques tailored to medical imaging datasets. These operations improve data consistency and reliability while supporting stronger predictive performance in the downstream learning models.

3.3. Local Model Training

Each healthcare node independently trains a local deep learning model using its private dataset. The framework employs Convolutional Neural Networks (CNNs) for medical image classification and Long Short-Term Memory (LSTM) networks for sequential healthcare data analysis. Local model parameters are optimized through backpropagation and gradient descent. Since raw patient data never leaves the local infrastructure, the framework fundamentally eliminates the privacy risks associated with centralized data sharing.

The local loss function at healthcare node i is defined as:

$$L_i(w) = (1/n_i) \Sigma \ell(x, y; w)$$

where L_i(w) represents the local loss at node i, n_i denotes the number of training samples, and w represents the model parameters.

3.4. Encrypted Parameter Exchange

After local training, model parameters and gradients are encrypted prior to transmission to the federated aggregation server. The framework employs secure aggregation and encrypted parameter exchange mechanisms to prevent attackers from reconstructing sensitive patient information from shared model updates. The encryption process is expressed as:

$$E(W^I) = W^I + R^I$$

where $E(W^I)$ represents the encrypted model weights, W^I denotes the local model parameters, and R^I is a random masking vector used for secure encryption.

3.5. Federated Aggregation

Encrypted model parameters from all participating institutions are transmitted to the Federated Aggregation Server, which performs Weighted Federated Averaging (FedAvg) to produce a global healthcare intelligence model. The aggregation is expressed as:

$$W\_global = \Sigma_i (n\_i / n) \cdot W^I$$

where W_global represents the global model weights, K is the number of healthcare nodes, n_i is the sample count at node i, and n is the total training samples across all institutions. The server combines encrypted updates without accessing raw patient records, preserving privacy throughout the aggregation process.

3.6. Blockchain Validation and PoS-BFT Consensus

Each local model update is submitted to the blockchain network for validation. Smart contracts automatically authenticate healthcare organizations, verify digital signatures, detect malicious updates, and prevent model poisoning attacks. The blockchain layer maintains a tamper-resistant distributed ledger of all verified model transactions. The verification function is expressed as:

$$V^I = H (E (W^I || T^I))$$

where V^I is the verification hash, H denotes the cryptographic hash function, $E(W^I)$ is the encrypted model update, and T^I is the transaction timestamp.

Unlike computationally expensive Proof-of-Work (PoW) systems, the proposed framework employs a lightweight PoS-BFT consensus mechanism that substantially reduces blockchain latency, computational complexity, and energy consumption while maintaining strong resilience against Byzantine failures and malicious participants. Validator nodes collaboratively verify blockchain transactions and ensure decentralized trust across the healthcare network. Figure 3 illustrates the blockchain validation and PoS-BFT consensus mechanism.

Following successful validation and consensus confirmation, the updated global model is redistributed to all participating institutions for the next federated learning round. This iterative process continues until the global model achieves optimal convergence. The final trained model is then deployed for clinical applications including disease diagnosis, medical image classification, remote patient monitoring, personalized treatment recommendation, and clinical decision support.

IV. EPP-BCFL ALGORITHM

Algorithm 1 outlines the complete workflow of the proposed EPP-BCFL framework:

Algorithm 1: Secure Federated Blockchain Learning (EPP-BCFL)

Input: Local healthcare datasets; Blockchain validator nodes

Output: Secure global healthcare model W_global

1. Initialize global model parameters W_g
2. Distribute W_g to all participating healthcare nodes
3. FOR each federated communication round t = 1, 2, .., T:
4. FOR each healthcare node i in parallel:
5. Train local model using private dataset D_i
6. Compute local loss L_i(W_i)
7. Encrypt local parameters: E(W_i) = W_i + R_i
8. Transmit E(W_i) to aggregation server
9. END FOR
10. FOR each model update E(W_i):
11. Submit update to blockchain smart contract
12. Execute PoS-BFT consensus validation
13. Reject malicious updates; accept valid updates
14. END FOR
15. Aggregate validated updates: W_global = Σ(n_i/n) · W_i
16. Redistribute updated W_global to all nodes
17. Check convergence criterion
18. END FOR
19. Return final global model W_global

V. DATASET DETAILS AND EXPERIMENTAL SETUP

The proposed EPP-BCFL framework was experimentally evaluated using multiple real-world healthcare datasets to assess classification accuracy, privacy preservation, communication efficiency, scalability, and resilience against adversarial attacks. The experimental environment was designed to simulate collaborative healthcare intelligence among distributed hospitals and healthcare organizations operating in a federated setting.

5.1. Datasets

Three benchmark healthcare datasets were employed in this evaluation:

MIMIC-III Clinical Database:
This large-scale, de-identified clinical database encompasses ICU patient records including demographics, laboratory reports, disease diagnoses, physiological measurements, medication information, and clinical observations. Its breadth and realistic clinical coverage make it an established benchmark for evaluating privacy-preserving federated learning frameworks in healthcare.

Breast Cancer Wisconsin Diagnostic Dataset:
This dataset contains diagnostic features extracted from fine needle aspirate images, with binary classification targets distinguishing malignant from benign tumors across multiple cellular characteristics. It was employed to evaluate the disease classification capability of EPP-BCFL in distributed healthcare settings.

ChestX-ray14 Medical Imaging Dataset:
Comprising thousands of chest X-ray images annotated with multiple thoracic disease labels — including pneumonia, infiltration, edema, fibrosis, and cardiomegaly — this dataset was used to evaluate the framework's performance in distributed medical image classification. Its inherent data sensitivity makes it a particularly appropriate benchmark for privacy-preserving collaborative learning.

5.2. Data Distribution and Preprocessing

The healthcare datasets were partitioned among multiple edge nodes in a Non-Identically Distributed (Non-IID) manner to replicate realistic hospital conditions where patient distributions and disease characteristics differ significantly across institutions. Each participating institution maintained its local dataset independently, without sharing raw patient records.

Preprocessing steps applied across all datasets included missing value imputation, noise removal, Min-Max normalization for numerical attributes, and anonymization of patient-identifiable information in compliance with HIPAA and GDPR. For medical imaging data, additional preprocessing steps included image resizing, grayscale normalization, and augmentation techniques such as rotation and horizontal flipping to improve model generalization.

5.3. Implementation Environment

The EPP-BCFL framework was implemented using Python 3.10 with the TensorFlow deep learning library for local model training. Hyperledger Fabric served as the blockchain platform for decentralized model verification and smart contract execution. Experiments were conducted on a computing environment equipped with an NVIDIA RTX 3080 GPU, Intel Core i9 processor, and 32 GB RAM. Healthcare edge nodes were simulated as distributed hospital clients participating in federated learning over 100 communication rounds.

CNN models were used for medical image classification tasks, while LSTM models were applied for sequential EHR-based disease prediction. The Adam optimizer was employed with a learning rate of 0.001 and a batch size of 32.

VI. EXPERIMENTAL RESULTS AND PERFORMANCE ANALYSIS

The proposed EPP-BCFL framework was systematically compared against four baseline methods: a conventional Centralized CNN model, traditional Federated Learning (FL), Secure Federated Learning (Secure FL), and Blockchain-enabled Federated Learning (Blockchain-FL). Experiments were carried out across the three healthcare datasets using five simulated distributed hospital nodes operating under non-IID data conditions.

6.1. Classification Accuracy

Classification accuracy is a critical metric in healthcare intelligence, as predictive reliability directly influences clinical decision-making and patient outcomes. As shown in Table 1, the proposed EPP-BCFL framework achieved the highest accuracy of 95.2% across all evaluated approaches.

The centralized CNN model achieved 88.4% accuracy, reflecting the limitations of single-institution training. Traditional FL improved accuracy to 91.8% through collaborative learning. Secure FL and Blockchain-FL further raised accuracy to 93.1% and 94.2%, respectively, through progressively more robust security mechanisms. The proposed EPP-BCFL framework outperformed all baselines by combining encrypted parameter exchange, secure aggregation, smart contract validation, and PoS-BFT consensus, achieving 95.2% accuracy.

Table 1. Classification accuracy comparison across healthcare learning models.

Model	Accuracy (%)
CNN (Centralized)	88.4
Federated Learning (FL)	91.8
Secure FL	93.1
Blockchain-FL	94.2
Proposed EPP-BCFL	95.2

6.2. Precision, Recall, and F1-Score

In healthcare diagnostics, false positives and false negatives carry significant clinical consequences. Table 2 presents the precision, recall, and F1-score for each evaluated approach. The proposed EPP-BCFL framework achieved the highest values across all three metrics, with 95.0% precision, 95.3% recall, and a 95.1% F1-score, demonstrating well-balanced and reliable diagnostic performance.

Table 2. Precision, recall, and F1-score comparison of healthcare intelligence models.

Metric	CNN	FL	Secure FL	EPP-BCFL
Precision	87.2%	90.4%	92.1%	95.0%
Recall	86.8%	91.1%	92.8%	95.3%
F1-Score	87.0%	90.7%	92.4%	95.1%

6.3. Communication Overhead

Communication overhead is a practical bottleneck in federated learning, since model parameters must be repeatedly exchanged between institutions and the aggregation server. High communication costs reduce system scalability and increase round latency. The proposed EPP-BCFL framework addresses this through lightweight PoS-BFT consensus and efficient encrypted parameter aggregation, achieving approximately 25% lower communication overhead compared with conventional Blockchain-FL systems, as summarized in Table 3.

Table 3. Communication overhead comparison across federated healthcare learning frameworks.

Model	Communication Overhead
CNN	Very High
FL	High
Secure FL	Moderate
Blockchain-FL	Moderate
Proposed EPP-BCFL	Low

6.4. Blockchain Latency

Blockchain latency directly affects the responsiveness of collaborative AI systems in time-sensitive clinical environments. Conventional Proof-of-Work systems incur substantial validation delays and energy costs, which are impractical for real-time healthcare applications. The proposed PoS-BFT consensus mechanism reduces the average transaction validation time from 4.8 seconds (traditional Blockchain-FL) to 1.9 seconds, representing a 60% reduction, as shown in Table 4.

Table 4. Blockchain validation latency comparison.

Framework	Avg. Validation Latency
Traditional Blockchain-FL	4.8 seconds
Proposed EPP-BCFL	1.9 seconds

6.5. Adversarial Attack Resistance

Security and robustness are non-negotiable requirements in federated healthcare learning, where malicious participants may attempt to inject poisoned model updates to compromise the global model. The EPP-BCFL framework employs blockchain smart contracts, encrypted parameter exchange, and decentralized PoS-BFT consensus to detect and reject adversarial updates before they influence the aggregation process. As shown in Table 5, the proposed framework achieves Very High attack resistance, outperforming all baseline approaches.

Table 5. Adversarial attack resistance comparison of healthcare learning frameworks.

Framework	Attack Resistance
FL	Moderate
Secure FL	Moderate
Blockchain-FL	High
Proposed EPP-BCFL	Very High

6.6. Scalability and System Efficiency

Scalability is essential for real-world deployment in large healthcare systems involving hundreds of hospitals and thousands of IoMT devices. The EPP-BCFL framework's decentralized architecture avoids centralized bottlenecks, while the lightweight PoS-BFT mechanism minimizes per-round blockchain overhead and improves throughput. Experiments confirmed that the framework maintains stable accuracy and low latency as the number of participating healthcare nodes increases, demonstrating practical suitability for large-scale deployments.

6.7. Overall Performance Summary

Table 6 provides a consolidated comparison of the proposed EPP-BCFL framework against all baselines across the key performance dimensions. The proposed framework consistently achieves the best outcomes in accuracy, privacy, and attack resistance, validating the effectiveness of combining federated learning, blockchain smart contracts, encrypted parameter exchange, and PoS-BFT consensus.

Table 6. Overall performance comparison across healthcare learning models.

Method	Accuracy	Privacy	Attack Resistance
CNN	88.4%	Low	Low
FL	91.8%	High	Moderate
Secure FL	93.1%	High	Moderate
Blockchain-FL	94.2%	Very High	High
Proposed EPP-BCFL	95.2%	Very High	Excellent

VII. DISCUSSION

The experimental results collectively demonstrate that the proposed EPP-BCFL framework addresses the central limitations of existing federated learning approaches for healthcare intelligence. The integration of blockchain-based validation with encrypted parameter exchange and lightweight PoS-BFT consensus yields measurable improvements across all evaluated dimensions.

The accuracy improvement from 91.8% (traditional FL) to 95.2% (EPP-BCFL) can be attributed to two complementary factors. First, blockchain validation ensures that only integrity-verified model updates contribute to global aggregation, thereby reducing the noise introduced by faulty or malicious participants. Second, the secure aggregation mechanism enables more healthcare institutions to contribute without fear of data leakage, broadening the effective training distribution.

The 60% reduction in blockchain validation latency (from 4.8 s to 1.9 s) achieved through PoS-BFT consensus represents a practically meaningful improvement for real-time clinical decision support, where delayed responses can have direct patient safety implications. Similarly, the 25% reduction in communication overhead supports deployment in bandwidth-constrained hospital networks and resource-limited IoMT environments.

Compared with Alazab et al. [20] and Kumar et al. [19], the proposed framework achieves higher accuracy and lower latency while maintaining comparable privacy guarantees. This is primarily due to the lightweight PoS-BFT consensus, which replaces computationally expensive PoW mechanisms, and the use of both smart contract validation and encrypted aggregation in combination, rather than individually.

A limitation of the current study is the reliance on simulated distributed hospital environments rather than deployment across real federated hospital networks. Furthermore, the Non-IID data partitioning used in evaluation, while realistic, may not fully capture the heterogeneity observed in live clinical systems. Future work will address these limitations through pilot deployments and more fine-grained data heterogeneity modeling.

VIII. CONCLUSION

This paper has presented the Enhanced Privacy-Preserving Blockchain-Enabled Federated Learning (EPP-BCFL) framework, a secure and scalable architecture for intelligent healthcare analytics. By

integrating federated learning with blockchain technology and a lightweight PoS-BFT consensus mechanism, the framework overcomes the key limitations of existing approaches in terms of privacy preservation, model integrity, communication efficiency, and adversarial robustness.

Experimental evaluation across three real-world healthcare datasets — MIMIC-III, Breast Cancer Wisconsin, and ChestX-ray14 — demonstrated that the proposed framework achieves 95.2% classification accuracy with reduced communication overhead, significantly lower blockchain validation latency (1.9 s vs. 4.8 s), and excellent resistance against adversarial model poisoning attacks. Smart contract-based validation provides tamper-resistant and decentralized trust management, making the framework well-suited for multi-institutional healthcare deployments.

Future research will explore the incorporation of explainable AI (XAI) mechanisms to enhance clinical interpretability, the adoption of quantum-resistant cryptographic protocols to future-proof privacy guarantees, adaptive trust scoring for dynamic participant management, and edge intelligence integration for resource-constrained IoMT environments.

REFERENCES

[1] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. Y. Arcas, "Communication-efficient learning of deep networks from decentralized data," *Proc. 20th Int. Conf. Artif. Intell. Statist. (AISTATS)*, pp. 1273–1282, 2017.

[2] P. Kairouz, H. B. McMahan, B. Avent et al., "Advances and open problems in federated learning," Found. Trends Mach. Learn., vol. 14, no. 1–2, pp. 1–210, 2021.

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

[4] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," ACM Trans. Intell. Syst. Technol., vol. 10, no. 2, pp. 1–19, 2019.

[5] K. Bonawitz, V. Ivanov, B. Kreuter et al., "Practical secure aggregation for privacy-preserving machine learning," Proc. ACM SIGSAC Conf. Comput. Commun. Secur., pp. 1175–1191, 2017.

[6] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," Cryptography Mailing List, 2008.

[7] Z. Zheng, S. Xie, H. Dai, X. Chen, and H. Wang, "An overview of blockchain technology: Architecture, consensus, and future trends," IEEE Int. Congr. Big Data, pp. 557–564, 2017.

[8] X. Xu, I. Weber, and M. Staples, Architecture for Blockchain Applications. Springer, 2019.

[9] E. Androulaki, A. Barger, V. Bortnikov et al., "Hyperledger Fabric: A distributed operating system for permissioned blockchains," Proc. 13th EuroSys Conf., pp. 1–15, 2018.

[10] P. Singh and K. Chatterjee, "Cloud security challenges and blockchain solutions for healthcare systems," J. Netw. Comput. Appl., vol. 147, pp. 102–120, 2020.

[11] C. C. Agbo, Q. H. Mahmoud, and J. M. Eklund, "Blockchain technology in healthcare: A systematic review," Healthcare, vol. 7, no. 2, Art. no. 56, 2019.

[12] A. F. Hussein, N. ArunKumar, G. Ramirez-Gonzalez et al., "Federated learning for healthcare informatics," J. Healthcare Eng., vol. 2020, pp. 1–16, 2020.

[13] N. Rieke, J. Hancox, W. Li et al., "The future of digital health with federated learning," npj Digit. Med., vol. 3, no. 1, pp. 1–7, 2020.

[14] M. J. Sheller, G. A. Reina, B. Edwards et al., "Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data," Sci. Rep., vol. 10, no. 1, pp. 1–12, 2020.

[15] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang, "On the convergence of FedAvg on non-IID data," Proc. Int. Conf. Learn. Representations (ICLR), 2020.

[16] S. King and S. Nadal, "PPCoin: Peer-to-peer crypto-currency with proof-of-stake," Self-Published Paper, 2012.

[17] M. Castro and B. Liskov, "Practical Byzantine fault tolerance," Proc. 3rd Symp. Operating Syst. Design Implementation, pp. 173–186, 1999.

[18] Z. Zhou, X. Wang, L. Sun et al., "Blockchain-based federated learning for secure healthcare systems," IEEE Access, vol. 9, pp. 123–135, 2021.

[19] R. Kumar, N. Marchang, and R. Tripathi, "Distributed deep learning and blockchain-enabled secure healthcare framework," Future Gener. Comput. Syst., vol. 126, pp. 128–140, 2022.

[20] M. Alazab, S. Khan, S. S. R. Krishnan et al., "Secure and efficient blockchain-enabled federated learning for healthcare applications," IEEE Trans. *Ind. Informat.*, vol. 17, no. 12, pp. 1–10, 2021.