Grenze International Journal of Engineering and Technology, June Issue

CP-FST: A Cognitive Privacy-Aware Federated Split Transformer for Secure Collaborative Healthcare Intelligence Systems

Stafin JR Shebu[1] and Sophia S[2]
[1-2]Division of Artificial Intelligence and Machine Learning, Karunya Institute of Technology and Sciences, Coimbatore, India
Email: stafinjrshebu1801@gmail.com, sophia@karunya.edu

***Abstract*—Secure and privacy-preserving collaboration remains a critical challenge in modern digital healthcare, where sensitive patient data cannot be centrally shared due to regulatory and security constraints. Addressing this need, this paper presents a Cognitive Privacy-Aware Federated Split Transformer (CP-FST) framework for secure and collaborative healthcare intelligence across distributed medical institutions. The approach integrates federated learning, split learning, and transformer-based representation to enable joint model training without sharing raw patient data, thereby addressing privacy, regulatory, and security concerns in digital healthcare environments. A cognitive privacy controller dynamically determines optimal model partitioning to balance privacy risk, communication overhead, and computational efficiency, while differential privacy-aware self-attention and trust-based secure aggregation enhance robustness against inference and adversarial attacks. Comprehensive preprocessing of electronic health records ensures reliable, sequence-aware clinical representations suitable for distributed learning. Experimental evaluation demonstrates superior predictive accuracy, faster convergence, lower training loss, and improved scalability compared with conventional centralized and federated methods, achieving approximately 98.7% accuracy. Overall, CP-FST provides a secure, efficient, and privacy-preserving solution for real-world collaborative healthcare analytics and supports future extensions toward larger multi-institutional deployments and stronger privacy guarantees.**

***Keywords*— Federated Learning; Split Learning; Transformer-based Healthcare Analytics; Privacy-Preserving Machine Learning; Collaborative Healthcare Intelligence.**

I. INTRODUCTION

The healthcare systems have become very digitalized which has made it possible to collect patient data on a high volume based on electronic health records, medical imaging and physiological monitoring devices[1]. To take advantage of these data in order to use them in intelligent clinical decision support, it is essential that multiple institutions of health engage in collaborative learning. Nevertheless, collecting sensitive medical information in a central place draws significant privacy, security, and regulatory issues especially where stringent regulations exist like HIPAA and GDPR [2].

Federated learning has become an emerging promising paradigm, which enables several different institutions to collectively train machine learning models without exchanging raw data [3].

Grenze ID: 01.GIJET.12.2.621_26

Although it has its benefits, the current federated learning systems are susceptible to gradient inversion and membership inference attacks that may compromise privacy, and are highly ineffective in terms of communication overhead. Split learning also minimizes the exposure of data through splitting the model computation between clients and servers, but fixed split patterns can still contain sensitive information, and cannot adapt to non-homogeneous healthcare environments. In addition, the majority of existing methods are based on traditional neural designs that are not capable of capturing multimodal healthcare data complex long-range dependencies [4], [5], [6]. To overcome these shortcomings, this paper presents a Cognitive Privacy-Conscious Federated Split Transformer (CP-FST) model of collaborative healthcare intelligence security. The major contributions of the present research are:

- Combines representation learning based on transformers with federated split learning to improve the protection of privacy, as well as maintaining the performance of the models.
- A privacy-conscious tool is dynamically used to regulate partitioning of a model to achieve a compromise between privacy risk and efficiency in communication.
- The mechanism of self- attention has been implemented with differential privacy directly built into them to reduce inference attacks.
- Strong resistance of malicious participants is also improved by secure and trust-aware aggregation.
- Demonstrates better privacy-utility trade-offs and robustness, therefore, applicable to practical collaborative healthcare implementations

Section I introduction covers the background and contributions, section II covers the literature review, section III covers the proposed framework and section IV result and discussion and concluded in section V

II. Literature Review

Recent studies underscore the growing importance of privacy-preserving collaborative learning in healthcare analytics due to sensitive patient data, regulatory requirements, and the widespread use of federated learning approaches. Zhang et al. [7] proposed SecureHealth, a multi-cloud federated architecture combining RNS-CKKS homomorphic encryption, adaptive differential privacy, and secure multi-party computation to get near-centralized accuracy with moderate computation overhead. A decentralized collaborative learning framework called DeCaPH was proposed by Fang et al. [8] it improves the privacy and predictive results with multi-hospital datasets, but its limited performance in scalability and convergence is caused by the absence of central coordination. Kuo et al. [9] proposed D-CLEF which supports horizontally and vertically partitioned healthcare data with performance almost centralized but at a high synchronization cost. Ahmed et al. [10] proposed a blockchain-assisted federated learning framework for healthcare monitoring systems to enhance the collaborative intelligence and communication efficiency for AI-enabled medical environments. Benabderrahmane et.al [11] introduced MedGuard-FL, a combination of adaptive differential privacy, encryption, anomaly detection, and Byzantine-resilient aggregation between device and edge and cloud layers that was found to be highly accurate, low-latency, but complex to the system. The reviewed Healthcare Information Exchange systems by Merhej et al. [12] indicate the possible synergies between blockchain and AI, but with a lack of empirical validation Zhang et al. [13] that SplitAVG would be an improvement of generalization based on network splitting and feature-map concatenation, with an increase in model complexity. Thakur et al. in [14] proposed a knowledge abstraction and filtering based federated learning framework for heterogeneous healthcare data to improve the collaborative model learning efficiency and reduce the communication complexity In general, existing approaches suffer from trade-offs between privacy preservation, scalability, computational efficiency, communication overhead, and real-world healthcare deployment.

A. Research Gap

Although existing studies demonstrate strong privacy guarantees and acceptable predictive performance, several gaps remain. Current federated and decentralized healthcare learning frameworks such as SecureHealth and MedGuard-FL rely heavily on cryptographic and differential privacy mechanisms, leading to increased computational and communication overhead that limits scalability and real-time deployment [7], [11]. Frameworks like DeCaPH and D-CLEF address decentralization and heterogeneous data but lack adaptive trust-aware weighting and efficient synchronization strategies [8], [9]. Moreover, blockchain- and ETL-based solutions focus primarily on secure data exchange rather than integrated learning intelligence. Finally, organizational and strategic intelligence perspectives are not yet technically unified with privacy-preserving collaborative learning, indicating the need for a lightweight, trust-aware, and scalable healthcare analytics framework.

III. RESEARCH FRAMEWORK

The proposed CP-FST framework integrates federated split learning with transformer-based representation modeling to enable privacy-preserving collaborative healthcare analytics across distributed institutions. Local feature extraction is performed at edge clients, while high-level representations are aggregated at a central server, ensuring data privacy. Privacy preservation is achieved through encrypted intermediate feature transmission, differential privacy-aware attention mechanisms, and trust-aware secure aggregation. The framework achieves scalable learning with reduced communication overhead and robust predictive performance across institutions, which is illustrated in Fig.1.

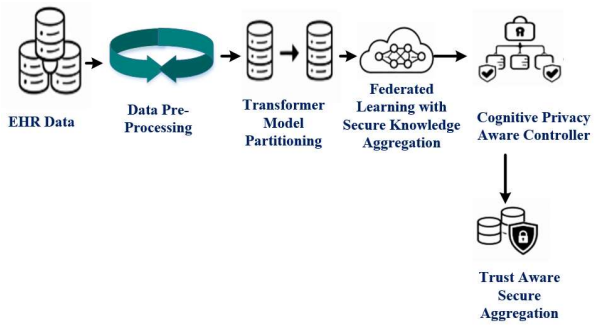

Fig. 1. Proposed Framework

A. Data Collection

The Electronic Health Record (EHR) data includes clinical data as a structured set of information that includes patient demographics, hospital admissions, and unit-level treatment outcomes. Patient-level identifiers are used to identify each patient and demographic attributes are age, gender, and ethnicity. The dataset captures detailed data of the hospital which includes hospital and ward identification, admission diagnosis, admission source, admission and discharge time, and discharge status and location. Besides these, fine unit-level data reflect the type of care unit, admission source, visit frequency, type of stay, body weight measurements, as well as unit discharge outcomes. [15].

B. Data Preprocessing

EHR preprocessing is performed locally at each institution to handle heterogeneous clinical distributions, missing values, temporal inconsistencies, and noisy physiological measurements before federated transformer training. The entire preprocessing procedure is done at the individual healthcare setup to maintain patient confidentiality.

Data Cleaning & Missing Value Handling

Data cleaning enhances the reliability and overall quality of the dataset by identifying and eliminating inaccurate, inconsistent, and erroneous record. Assume that the dataset is represented as eqn. (1)

$$D = \{x_1, x_2, x_3, \ldots, x_n\} \tag{1}$$

where $x_i \in R^m$ represents the feature vector of the i^{th} patient record.Every record is verified by a rule-based data consistency, including maintaining a valid age range, positive physiological values, and permissible categorical value. Records that violate clinical constraints are marked and removed as presented in eqn.(2)

$$x_{ij} = \frac{1}{N} \sum_{k=1}^{N} x_{kj} \tag{2}$$

Adaptive Missing Value Estimation

Group-aware imputation is imposed instead of global mean imputation. The replacement of a missing feature value x_ij is calculated using other patients that share demographic characteristics

$$x_{ij} = \frac{1}{|S|} \sum_{k\in}^{i} x_{kJ} \tag{3}$$

S denotes the set of records with the same gender and age bracket. This maintains clinical relevance and decreases bias caused by homogenous averaging.

Robust Outlier Suppression

The Interquartile Range (IQR) method is used to identify the outliers as opposed to the standard deviation as given in eqn.(4)

$$x_{ij} \notin [Q_1 - \alpha.IQR, Q_3 + \alpha.IQR] \quad (4)$$

Where Q_1 and Q_3 represent the first and third quartile, IQR=$Q_3 - Q_1$ and α=alpha is a scaling factor. Boundary limits are clipped instead of being removed to prevent the loss of data as a result of having extreme values.

Temporal Feature Transformation

Admission and discharge time are converted to cyclic representations in order to reflect periodic patterns of the hospital .this is expressed as eqn.(5)

$$t_{sin} = \sin\left(\frac{2\pi}{24}\right), t_{cos} = \cos\left(\frac{2\pi}{24}\right), \quad (5)$$

This encoding would maintain the continuity of time and enhance learning of transformer attention

Feature Scaling Using Robust Normalization

In order to reduce the impact of distorted distributions, normalization based on median is used this is expressed in eqn.(6)

$$x^1 = \frac{x - median(x)}{IQR(x)} \quad (6)$$

This enhances stability when there is heterogeneous data distributions of institutional data .The categorical attributes are transformed into dense embedding vectors, whereas events that occur to a patient are organized into visit-long sequences of fixed length.

C. Cognitive Privacy-Aware Federated Split Transformer (CP-FST) Architecture

The Cognitive Privacy-Aware Federated Split Transformer (CP-FST) system is meant to provide confidential and privacy-preserving joint learning using spread out healthcare information. In CP-FST, the transformer network is divided between the clients (e.g., hospitals, clinics, wearable devices) and a central server. The raw healthcare data do not move out of the client devices. Rather, the range of features sent is intermediary, which greatly limits the privacy loss and makes it highly competent in learning.

Transformer Model Partitioning (Split Learning Mechanism)

The core transformer model is subdivided in two logical parts:

Client-Side Transformer Module: Clients generate a packet format that is transmitted to the server (client) and returned to the client as a response packet.

The client-side module processes raw healthcare data locally. It consists of Input embedding layers, Positional encoding, and initial transformer encoder block.

Let the raw healthcare input at client i be represented as eqn.(7)

$$X_i \in R^{t \times f} \quad (7)$$

where T represents the length of sequences (e.g. time steps in physiological signals) and F is the dimension of features. The client calculates a latent representation, this is expressed as eqn. (8)

$$H_i = f_c(X_i; \theta_c) \quad (8)$$

where f_c denotes the client-side transformer function and θ represents client-side parameters.

Only the intermediate representation H_i is transmitted to the server, ensuring that sensitive raw patient data remain local.

Server-Side Transformer Module: The server-side module processes intermediate features received from clients to perform global learning and prediction. The server takes feature embeddings of a group of clients and engages in global representation learning as presented in eqn. (9)

$$y_i = f_s(H_i; \theta_s) \quad (9)$$

where, $f_s$is the server-side transformer map and thesis the server-side parameters. This allows for efficient global representation learning with lower parameter transmission overhead compared to traditional federated learning architectures.

Federated Learning with Secure Knowledge Aggregation

CP-FST has a federated training scheme, in which the updates to the models are learnt in a collaborative manner without exchanging raw data. Encrypted gradients or split-layer updates are only exchanged after local forward and backward propagation. An aggregation mechanism that is secure is used to get updates combined in eqn.(10)

$$\theta_s^{(t+1)} = \sum_{i=1}^{N} w_i \theta_{s,i}^{(t)} \qquad (10)$$

Where N is the number of participating clients w_i is a trust-aware weighting factor $\theta_{s,i}^{(t)}$ denotes server-side updates from client *i.* This Trust-aware aggregation improves robustness against adversarial updates, model poisoning attacks, and statistically inconsistent client behavior during collaborative optimization.

Cognitive Privacy-Aware Controller

A key contribution of CP-FST is the adaptive privacy-aware controller that dynamically optimizes transformer split locations according to privacy risk, communication cost, and client computational capacity. The controller considers: The split point optimization is formulated by jointly minimizing the privacy risk level (P_r) and communication cost (C_c) while respecting the client, $C_c$computational capacity. The objective is to determine an optimal split point that balances privacy preservation, efficient data transmission, and feasible on-device computation as presented in eqn. (11)

$$\min_s \alpha P_r(s) + \beta C_c(s) + \gamma C_p(s) \qquad (11)$$

where α, β, γ are adaptive weighting coefficients and *s* denotes the split layer index.

iv)Differential Privacy-Aware Self-Attention Mechanism

CP-FST implements differential privacy (DP) as a direct part of the transformer self-attention mechanism to prevent inference and reconstruction attacks. The score of attention in normal attention is calculated as following eqn. (12)

$$A = softmax(\frac{QK^T}{\sqrt{d_k}} \qquad (12)$$

CP-FST injects calibrated Gaussian noise as presented in eqn. (13)

$$A' = softmax(\frac{QK^T}{\sqrt{d_k}} + N(0,\sigma^2)) \qquad (13)$$

where σ is determined by the privacy budget ε.This controlled noise addition mechanism enhances the security of the CP-FST framework by effectively limiting the amount of sensitive information that can be inferred from intermediate representations or attention scores.

Trust-Aware Secure Aggregation

To further strengthen security, CP-FST incorporates a trust evaluation mechanism for participating clients. Each client is assigned a trust score computed based on Update consistency, Historical reliability, Deviation from global model behaviour. The aggregation T_i weight is adjusted in eqn. (14)

$$w_i = \frac{T_i}{\sum_{j=1}^{N} T_j} \qquad (14)$$

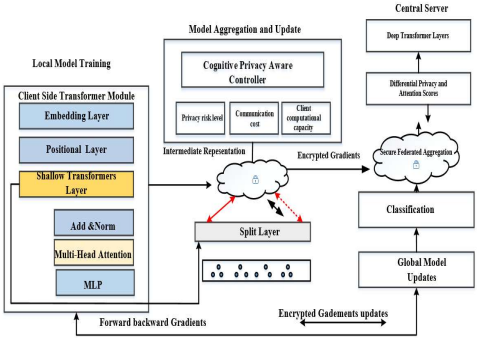

Fig. 2. Cognitive Privacy-Aware Federated Split Transformer (CP-FST) Framework

This approach effectively suppresses the impact of anomalous or malicious updates and enhances collaborative robustness. The CP-FST training procedure begins with clients performing local forward propagation up to the dynamically determined split layer, after which only intermediate embeddings are transmitted to the server Fig. 2 illustrates the complete CP-FST workflow, including local transformer encoding, split-layer communication, privacy-aware attention learning, and trust-based federated aggregation across distributed healthcare institutions.

IV. Results and Discussions

The results show that the proposed Cognitive Privacy-Aware Federated Split Transformer (CP-FST) outperforms conventional centralized and federated models in both predictive accuracy and training efficiency. The framework achieves faster convergence, lower loss, and improved robustness to heterogeneous healthcare data across distributed clients. Its split-transformer design reduces computational and communication overhead by distributing processing between edge devices and the server. At the same time, integrated privacy-preserving mechanisms effectively protect sensitive patient information without sacrificing model performance. Overall, the system demonstrates secure, scalable, and reliable collaborative healthcare intelligence

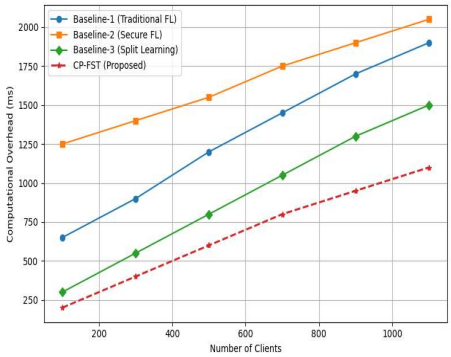

Fig. 3. Computational Overhead

Fig.3. used to compare the computational overhead of various learning procedures with increase in the number of clients.CP-FST is the least computationally expensive of all client sizes, and is also more scalable and efficient.Conventional and safe federated learning approaches incur much more overhead particularly when there are more clients.The findings, in general, show that CP-FST minimizes the computational process but does not deteriorate the effect of collaborative learning.

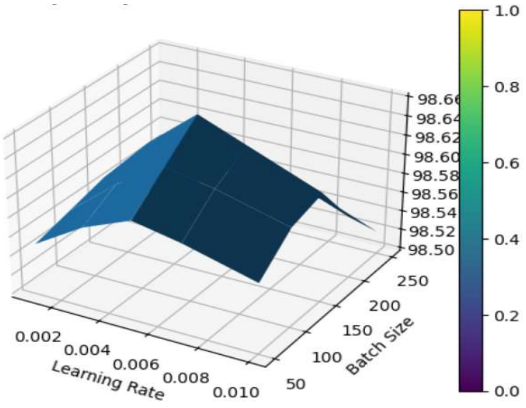

Fig. 4. Effect of Learning Rate and Batch Size on Model Accuracy

Fig.4. illustrates the privacy-utility trade off of the CP-FST framework, which demonstrates how the accuracy varies with varying the learning rates and batch sizes with privacy constraint. CP-FST shows very high accuracy (more than 98.5) which means that it is very robust to changes in hyperparameters. Optimal performance is achieved with medium learning rates and medium batch sizes though extreme settings slightly decrease accuracy

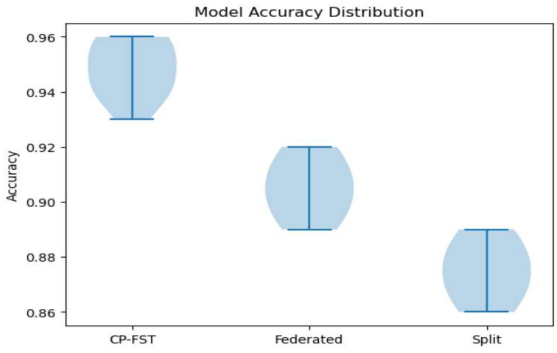

Fig. 5. Accuracy distribution

Fig.5. demonstrates the accuracy distribution of various models (CP-FST, Federated, and Split) in terms of various runs or clients.The broader parts reveal the concentrations of the majority of the accuracy values.CP-FST is more and better accurate than the other models.It shows that CP-FST is more stable and achieves better performance in federated learning.

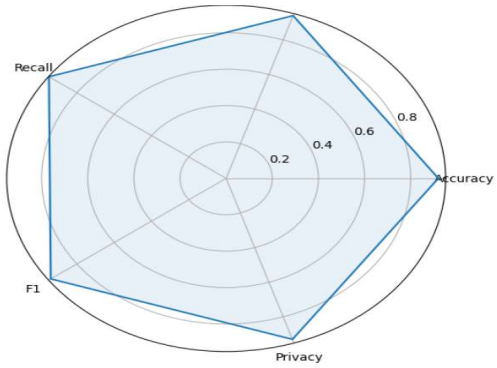

Fig. 6. Multi-performance outcome

Fig.6.indicates the multi-metric results of the CP-FST model as per the Accuracy, Precision, Recall, F1-score, and Privacy. The evaluation metrics are represented by each axis, and the bigger the area, the more powerful is the overall model performance's-FST has almost a circular shape, which implies that the latter has an overall high score in all indicators without significant flaws.

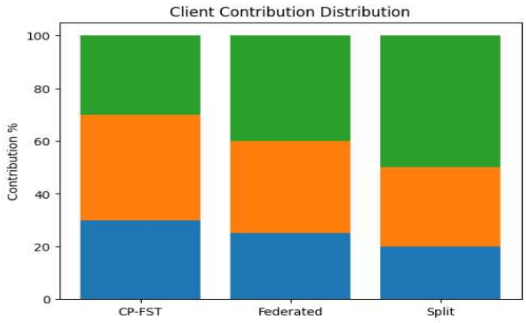

Fig. 7. Client Contribution Distributions

Fig.7. represents the contribution of the clients towards CP-FST, Federated and Split learning systems. CP-FST demonstrates a more balanced contribution between clients when compared to Federated and Split learning where the contribution of higher contributing clients are on the rise thus showing uneven contribution. CP-FST, in general, enhances more balanced and sustainable client engagement in collaborative training.

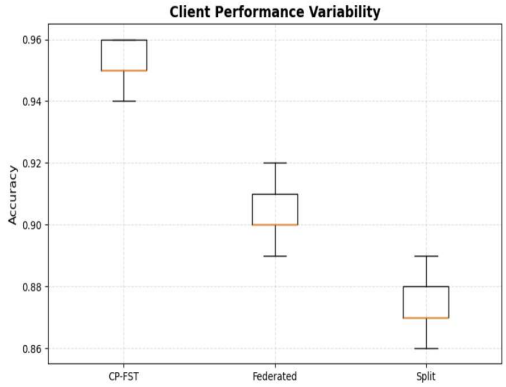

Fig. 8. Client Performance Variability Across CP-FST, Federated, and Split Learning Approaches

Fig.8. contrasts the variation of the client performance in CP-FST, Federated and Split learning. CP-FST boasts of the highest median accuracy and the least variability implying the same and consistent client performance. Federated and Split learning, on the other hand, possess lower accuracy and a broader spread, indicating more client inconsistency in performances.

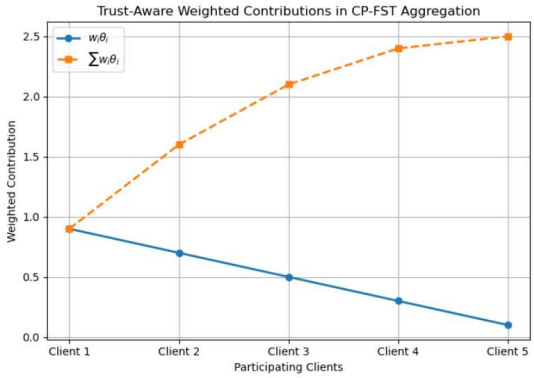

Fig 9. Trust-Aware Weighted Client Contributions in CP-FST Aggregation

Fig.9. shows that trust-dependent weighted contributions in CP-FST aggregation, where clients with higher trust score have higher weight in the model update. The cumulative aggregation is constant with smaller weights $w_i\theta_i$as the individual client weights decrease, and grows with larger weights w_i.

TABLE I: PERFORMANCE COMPARISON

Author / Year	Accuracy (%)
SplitAVG (Federated Split Learning) [13]	96.2
MedHE: [14]	96.8
MedGuard-FL [11]	95.6
DeCaPH Decentralized FL [8]	94.9
Proposed CP-FST	98.7

The table 1 compare the accuracy of the existing federated healthcare learning models with the proposed CP-FST framework. The proposed CP-FST has the best accuracy (98.7%), which means it has a better predictive ability. This has mainly been enhanced by the fact that it has transformer-based learning, adaptive split architecture, and privacy-sensitive secure aggregation. All in all, the findings verify that CP-FST is more efficient in learning and model performance than the current federated healthcare strategies.

V. CONCLUSION AND FUTURE WORK

The present research is based on a Cognitive Privacy-Aware Federated Split Transformer (CP-FST) architecture proposed in this paper to make healthcare intelligence secure and collaborative across the distributed medical

facilities. The combination of federated learning, split learning, and transformer based models makes the framework highly predictive without allowing any violation of data privacy. Patient data are stored in sensitive local edge computers, and only encrypted parameters or intermediate representations are exchanged with the central server, minimizing the risk of data leaks and regulatory issues. The results of the experiment prove the high quality of classification, reduced convergence time, reduced training loss, and improved generalization of centralized and standard federated strategies. The split computation strategy decreases the communication overhead and computational load and thus, the system is appropriate to resource-constrained environments and can be scaled across multiple hospitals even in the presence of data heterogeneity. Despite these difficulties like network latency and synchronization complexity, further development will involve how to make communication efficient by adapting clients selection and designs of light weight transformers. Other privacy controls such as differential privacy and secure multiparty computation and testing on a large-scale real-world data will make it even more reliable and applicable.

REFERENCES

[1] Z. Wu, J. Hou, Y. Diao, and B. He, "Federated transformer: Multi-party vertical federated learning on practical fuzzily linked data," Advances in Neural Information Processing Systems, vol. 37, pp. 45791–45818, 2024.

[2] V. Krishnaprasath, V. Pamisetty, V. Sharma, M. Nayak, N. Baalakumar, and S. Aravindh, "Federated Learning Based Artificial Intelligence Systems with Blockchain Security for Global Healthcare Collaboration and Patient Centric Data Privacy," in International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 2024), Atlantis Press, 2025, pp. 1277–1290.

[3] A. Raza, Secure and privacy-preserving federated learning with explainable artificial intelligence for smart healthcare system. University of Kent (United Kingdom), 2023.

[4] S. A. Alzakari, A. Sarkar, M. Z. Khan, and A. A. Alhussan, "Converging technologies for health prediction and intrusion detection in internet of healthcare things with matrix-valued neural coordinated federated intelligence," IEEE Access, vol. 12, pp. 99469–99498, 2024.

[5] G. C. Amaizu, A. M. V. V. Sai, S. Bwardwaj, D.-S. Kim, M. Siddula, and Y. Li, "FedViTBloc: secure and privacy-enhanced medical image analysis with federated vision transformer and blockchain," High-Confidence Computing, p. 100302, 2025.

[6] X. Ma, J. Hu, S. Liang, and Y. Wu, "Federated Learning and Resource-Aware Graph Neural Network for Intrusion Detection in 6G-IoT Driven Healthcare System," IEEE Internet of Things Journal, 2025.

[7] H. Zhang, E. Feng, and H. Lian, "A Privacy-Preserving Federated Learning Framework for Healthcare Big Data Analytics in Multi-Cloud Environments," Spectrum of Research, vol. 4, no. 1, 2024.

[8] C. Fang et al., "Decentralised, collaborative, and privacy-preserving machine learning for multi-hospital data," EBioMedicine, vol. 101, 2024.

[9] T.-T. Kuo, R. A. Gabriel, J. Koola, R. T. Schooley, and L. Ohno-Machado, "Distributed cross-learning for equitable federated models-privacy-preserving prediction on data from five California hospitals," Nature Communications, vol. 16, no. 1, p. 1371, 2025.

[10] J.-Y. Ma and T.-W. Kang, "Digital Intelligence and Decision Optimization in Healthcare Supply Chain Management: The Mediating Roles of Innovation Capability and Supply Chain Resilience," Sustainability, vol. 17, no. 15, p. 6706, 2025.

[11] F. Benabderrahmane, E. Kerkouche, and N. Bouchemal, "Risk-Aware Privacy-Preserving Federated Learning for Remote Patient Monitoring: A Multi-Layer Adaptive Security Framework," Applied Sciences, vol. 16, no. 1, p. 29, Dec. 2025, doi: 10.3390/app16010029.

[12] J. Merhej, H. Harb, A. Abouaissa, and L. Idoumghar, "Toward a New Era of Smart and Secure Healthcare Information Exchange Systems: Combining Blockchain and Artificial Intelligence," Applied Sciences, vol. 14, no. 19, p. 8808, Sep. 2024, doi: 10.3390/app14198808.

[13] M. Zhang, L. Qu, P. Singh, J. Kalpathy-Cramer, and D. L. Rubin, "SplitAVG: A Heterogeneity-Aware Federated Deep Learning Method for Medical Imaging," IEEE J. Biomed. Health Inform., vol. 26, no. 9, pp. 4635–4644, Sep. 2022, doi: 10.1109/JBHI.2022.3185956.

[14] F. Yesmin, "MedHE: Communication-Efficient Privacy-Preserving Federated Learning with Adaptive Gradient Sparsification for Healthcare," arXiv preprint arXiv:2511.09043, 2025.

[15] "Electronic Health Record." Accessed: Feb. 09, 2026. [Online]. Available: https://www.kaggle.com/datasets/anuchhetry/electronic-health-record