Federated Learning with Homomorphic Encryption for Confidential Healthcare Data Sharing*

Domenico Lofù[1], Paolo Sorino[1], Carmelo Antonio Ardito[1], Tommaso Di Noia[1] and Eugenio Di Sciascio[1]

[1]Dept. of Electrical and Information Engineering (DEI), Politecnico di Bari, Bari, BA, 70125, Italy

Abstract

Recent advances in data-driven healthcare have highlighted the need for machine learning systems that can leverage sensitive clinical information without compromising patient privacy. In this context, Federated Learning (FL) has emerged as a promising paradigm, enabling collaborative model training across distributed institutions while keeping data local. However, strict privacy constraints and the risk of leakage through shared model updates remain significant challenges. This work presents a privacy-preserving FL architecture that integrates Homomorphic Encryption (HE) to secure end-to-end model updates. The proposed system, implemented using the CKKS scheme, enables the server to aggregate encrypted client parameters using a customised HEFedAvg strategy, preventing access to gradients or weights in plaintext. The proposed framework is validated on a distributed synthetic oncology dataset using a Multi-Layer Perceptron. Results show that encrypted FL achieves performance comparable to standard unencrypted training, reaching 92.40% accuracy and a weighted F1-score of 92.26%, while ensuring strong protection against inference attacks. Despite increased communication overhead due to ciphertext expansion, the approach remains efficient and fully compliant with clinical privacy requirements.

Keywords

Federated Learning, Homomorphic Encryption, Clinical Data Analysis

1. Introduction

The digital transformation of healthcare systems has and the adoption of Artificial Intelligence (AI) within, led to an unprecedented growth in the availability of clinical data, ranging from Electronic Health Records (EHRs) to diagnostic measurements and genomic information [1, 2, 3]. These data sources have the potential to fuel advanced Machine Learning (ML) models capable of supporting diagnosis, prognosis, and treatment personalisation [4, 5]. However, the sensitive nature of medical data, combined with stringent regulatory constraints such as the General Data Protection Regulation (GDPR)[1], makes its centralisation extremely challenging. As a result, a crucial research question arises: *how can healthcare institutions collaboratively train predictive models without compromising data confidentiality and sovereignty?*

Federated Learning (FL) has emerged as a promising response to this challenge. Indeed, FL enables multiple entities to train a shared ML model while keeping raw data localized at the source. Instead of transferring patient information, institutions communicate only model updates such as weights or gradients which are aggregated by a central server through federation algorithms such as FedAvg [6]. This paradigm is particularly suited to the healthcare domain, where privacy, data governance, and compliance requirements prohibit data pooling across institutions.

1st Workshop on Symbiotic AI for Cyber Social Security: Human–AI Collaboration for Social Resilience and Trustworthy Digital Ecosystems (SAI4CCS 2026), co-located with the 2nd Joint National Conference on Cybersecurity (ITASEC & SERICS 2026), February 09–13, 2026, Cagliari, Italy

domenico.lofu@poliba.it (D. Lofù); paolo.sorino@poliba.it (P. Sorino); carmelo.ardito@poliba.it (C. A. Ardito); tommaso.dinoia@poliba.it (T. D. Noia); eugenio.disciascio@poliba.it (E. D. Sciascio)

https://sisinflab.poliba.it/people/domenico-lofu/ (D. Lofù); https://sisinflab.poliba.it/people/paolo-sorino/ (P. Sorino); https://sisinflab.poliba.it/people/carmelo-antonio-ardito/ (C. A. Ardito); https://sisinflab.poliba.it/people/tommaso-di-noia/ (T. D. Noia); https://sisinflab.poliba.it/people/eugenio-di-sciascio/ (E. D. Sciascio)

0000-0001-6413-9886 (D. Lofù); 0000-0002-9081-2648 (P. Sorino); 0000-0001-8993-9855 (C. A. Ardito); 0000-0002-0939-5462 (T. D. Noia); 0000-0002-5484-9945 (E. D. Sciascio)

© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

[1]https://eur-lex.europa.eu/eli/reg/2016/679/

CEUR-WS.org/Vol-4329/paper8.pdf

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073

published 2026-11-14

Despite its advantages, FL does not inherently guarantee complete privacy protection. Indeed, a series of attacks during training can leak sensitive information through model inversion or inference attacks, enabling the potential reconstruction of private clinical attributes [7, 8]. The risk is amplified in tabular medical datasets, where features often carry direct clinical meaning, *e.g.* tumour size, comorbidities, or genetic markers, making them more vulnerable to re-identification. Moreover, FL in healthcare faces additional challenges, including statistical heterogeneity across hospitals, differences in coding standards, non-Independent and Identically Distributed (IID) data distributions, and variability in infrastructure capabilities.

To mitigate these issues, advanced privacy-preserving mechanisms such as Differential Privacy [9], Secure Multi-Party Computation [10], and Homomorphic Encryption (HE) [11] are proposed. Among these, HE stands out as a mathematically rigorous solution enabling computations directly on encrypted data. In the FL context, HE allows the server to aggregate encrypted model updates without ever accessing them in plaintext, thereby ensuring end-to-end confidentiality even in the presence of an untrusted or potentially compromised aggregator. In particular, lattice-based schemes such as CKKS support efficient approximate arithmetic on floating-point vectors, making them suitable for neural network gradients and weights.

This study proposes a privacy-preserving FL architecture for clinical data, integrating the CKKS homomorphic encryption scheme into the FL workflow. The contributions of this work are threefold:

- A fully operational federated architecture integrating HE for end-to-end protection of model updates.
- A complete experimental validation demonstrating that encrypted FL can match the performance of non-encrypted training, achieving over 92% accuracy while maintaining complete data confidentiality.
- An analysis of computational and communication overhead, highlighting trade-offs and demonstrating the feasibility of HE-based FL in realistic healthcare scenarios.

Overall, the proposed approach demonstrates that FL and HE can be effectively combined to create secure, scalable, and compliant infrastructures for multi-institutional clinical collaboration. This supports the development of ML models that respect patient privacy while enabling the full potential of distributed healthcare data.

2. Related Work

Current literature has not yet fully established how effectively FL can ensure reliable performance and interoperability in downstream clinical tasks an open challenge that the present work aims to address. FL has been shown to achieve accuracy comparable to centralized learning while keeping data on premises, thereby meeting legal and ethical constraints on clinical data exchange [12]. In such settings, institutions collaboratively train a shared model by communicating only model updates, which are subsequently aggregated by a central server.

A wide range of studies has applied FL to imaging tasks such as brain tumour segmentation, breast density classification, and cancer detection from radiological or histopathological data [12, 13, 14, 15, 16]. These works indicate that FL supports heterogeneous multi-centre datasets while maintaining strict data locality, thereby improving model generalisation across clinical acquisition settings. More recent frameworks also integrate explainability and robustness mechanisms into FL pipelines for oncological imaging [17].

In addition to its application in imaging, FL has been widely explored for structured clinical data and EHRs. Numerous studies propose federated models for risk prediction, outcome modelling, and disease classification based on tabular representations of routine clinical information [18, 19]. This line of research identifies persistent challenges in working with EHR-like data, including pronounced class imbalance, institutional variability in coding practices, and the need for comprehensive preprocessing and feature harmonisation. Systematic reviews further highlight the increasing relevance of FL across

a broad range of medical specialties, including cardiology, oncology, intensive care, and population health management [20, 21].

In parallel, a growing number of studies focus on privacy-preserving FL in healthcare. This literature examines threats such as gradient inversion, membership inference, and model poisoning, and surveys mitigation strategies based on Differential Privacy (DP), Secure Multi-Party Computation (SMPC), and Homomorphic Encryption (HE) [22, 23, 24]. DP provides formal privacy guarantees via noise injection, albeit with possible accuracy degradation. SMPC offers strong cryptographic protection but often incurs high computational and communication overhead. HE specially CKKS-based schemes offers a balanced approach by enabling computation directly on encrypted floating-point data while preserving convergence properties [25].

Despite the increasing attention to privacy-preserving FL, fully homomorphically encrypted FL architectures for structured clinical data remain limited. Existing HE-based FL approaches in the medical domain predominantly focus on imaging tasks or rely on simplified threat models and non-representative datasets. Motivated by this gap, the present work introduces the first federated architecture that unifies the Flower framework with a CKKS-based HE backend and a bespoke *HEFedAvg* aggregation strategy, specifically engineered for secure multi-class oncology classification using tabular clinical features. Our results demonstrate that high predictive performance can be achieved while preserving end-to-end cryptographic protection of model updates, thereby offering a practical and deployable solution for advancing privacy-preserving clinical data sharing.

3. Proposed Architecture

This section introduces the proposed privacy-preserving FL architecture for secure training of clinical predictive models across multiple healthcare institutions. The architecture integrates HE with a fully operational FL workflow, ensuring that model updates remain encrypted throughout their entire life cycle, from local training to global aggregation, thus preventing any leakage of sensitive medical information. The proposed architecture extends the classical FL paradigm by incorporating cryptographic protection, robust data preprocessing, and tailored aggregation strategies to address the unique constraints of healthcare environments.

As shown in Figure 2, the architecture adopts a hub-and-spoke FL architecture [26], where a aggregator coordinates a federation of hospital clients without ever accessing raw patient data. Each hospital (node) trains a local model on its own EHR, while only encrypted model updates are transmitted to the server (aggregator). The system is implemented using the Flower framework[2] for FL orchestration and TenSEAL[3] for homomorphic encryption, following the CKKS scheme. This design balances computational feasibility with compliance with data protection regulations (*e.g.* GDPR).

3.1. Data Preprocessing Module

Given the heterogeneity of clinical datasets, the framework incorporates a rigorous local preprocessing pipeline executed exclusively on each client node. The pipeline includes the following steps: (i) stratified partitioning to preserve class distributions during cross-client data allocation, (ii) *Data Cleaning*, removing incomplete or invalid clinical records, (iii) *Feature Encoding*, combining One-Hot Encoding for categorical variables and ordinal encodings for structured clinical attributes and, (iv) *Feature Scaling* via Z-score normalization to ensure stable optimization across clients.

These steps guarantee that model updates produced by different hospitals are statistically coherent and suitable for aggregation, despite the intrinsic Non-IID nature of distributed clinical data.

[2] https://flower.ai
[3] https://github.com/OpenMined/TenSEAL

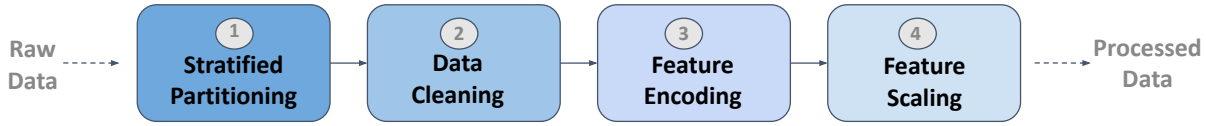


Figure 1: Representation of the different modules included within the data processing pipeline.

3.2. Local Training Module

Each client trains a Multi Layer Perceptron (MLP) optimized on its own clinical data. The MLP architecture consists of two fully connected hidden layers (64 and 32 units) with ReLU activations and an output layer predicting tumor classes (lung, liver, stomach). The scikit-learn[4] implementation supports incremental training via $partial_fit$, enabling local updates over multiple mini-batches.

Moreover, the federated training is supported by the following choices: (i) warm-start training, preserving optimizer state across rounds, (ii) Small local epoch count, limiting client drift under Non-IID data and, (ii) Class consistency mechanisms, injecting dummy samples when local batches lack some target labels, thus ensuring dimensional consistency across updates. This ensures compatibility with secure aggregation and provides reliable convergence properties.

3.3. Security Module

The security module relies on the CKKS homomorphic encryption scheme, suitable for neural parameters represented as floating-point values. The aggregator generates a cryptographic context containing the Public key (pk), the Private key (sk), and the Evaluation key (evk). Only the public context is distributed to clients, enabling them to serialize local model parameters into vectors, encrypt parameters using CKKS and, transmit ciphertexts as byte arrays.

Encrypted updates remain protected even if communication channels or the central server are compromised. The ciphertext expansion, determined by the polynomial modulus degree (8192) and coefficient modulus chain, introduces additional communication cost but does not impair model correctness.

3.4. Secure Aggregation Module

To support encrypted parameters, the FedAvg [27] algorithm was extended to HEFedAvg [28], a custom aggregation strategy that sums encrypted model updates without decryption. For each training round the server receives ciphertext vectors from all clients. Each vector is deserialized into a CKKS ciphertext. Subsequently, an weighted aggregation is performed directly in the encrypted domain (homomorphic addition and scalar multiplication). Only after aggregation, the server decrypts the combined ciphertext to obtain the updated global model. This approach preserves the mathematical properties of FedAvg while ensuring that the server never observes individual client updates in plaintext, thus eliminating major attack vectors such as gradient inversion.

Figure 2 summarize the complete training workflow executed on our system. It can be summarized as follows:

1. **Initialization Phase:** The server generates the CKKS context and broadcasts the public context along with the initial global model.
2. **Local Update Phase:** Each client trains the MLP for a fixed number of epochs using its private EHR partition.
3. **Encryption Phase:** The updated parameters are encrypted with the public CKKS key.
4. **Transmission Phase:** Encrypted updates are sent to the server through secure Flower communication channels.
5. **Homomorphic Aggregation Phase:** HEFedAvg aggregates encrypted updates from all participating clients.

[4]https://scikit-learn.org/stable/

6. **Decryption and Global Update Phase:** The server decrypts and normalizes the aggregated model.
7. **Iteration Phase:** The updated global model is redistributed for the next round.

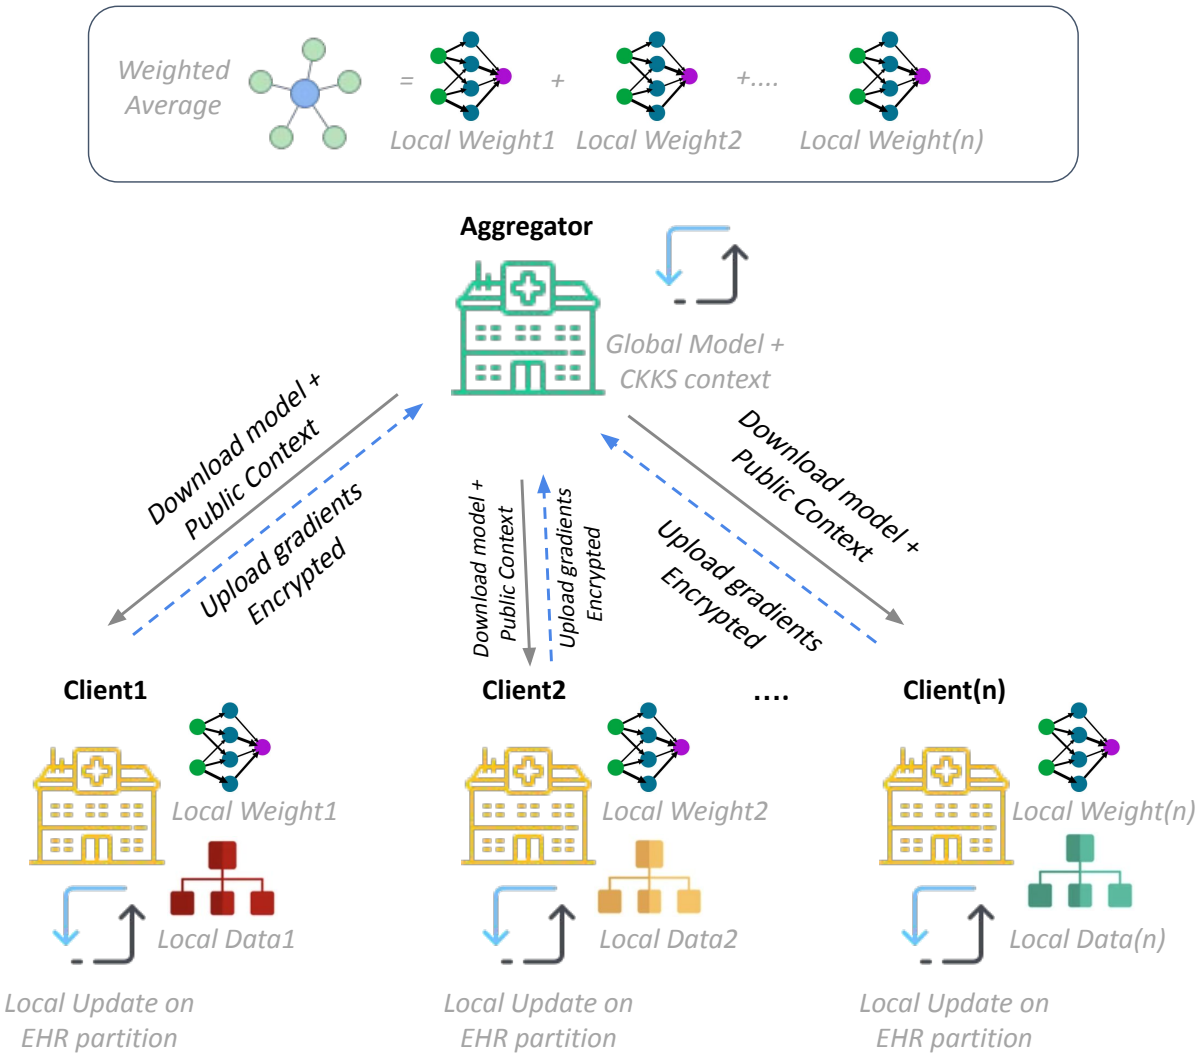


Figure 2: Representation of the federated training pipeline. The server initializes the CKKS context and global model; clients perform local MLP updates on private EHR partitions, encrypt parameters using the public CKKS key, and transmit ciphertexts via secure Flower channels. HEFedAvg performs homomorphic weighted aggregation, after which the server decrypts, normalizes, and redistributes the updated global model for the next round.

This cycle is repeated until convergence, while metrics such as accuracy, precision, and F1-score are monitored at the server to understand how well the model is performing.

4. Experimental Settings

This section describes the experimental setup adopted to evaluate the proposed privacy-preserving FL architecture based on HE. This phase consists of four key sub-phases: *Dataset Analysis*, *Preprocessing Pipeline*, *Federated Learning Configuration* and *Homomorphic Encryption Setup*.

4.1. Dataset Analysis and Preprocessing Pipeline

Experiments were conducted using a standardized synthetic clinical dataset designed to emulate EHRs of oncology patients. The dataset comprises $10,000$ samples and includes 20 features. Prior to training, an Exploratory Data Analysis (EDA) was performed to inspect class distribution, missing values, and correlations between numerical variables. The target variable $TumorType$ consists of three balanced classes: $Lung$, $Liver$, and $Stomach$.

A three-stage preprocessing workflow was adopted to ensure consistency of the federated training process across heterogeneous clients:

1. **Data Cleaning:** Non-informative attributes (patient identifiers, administrative timestamps) were removed. Samples with missing values in essential clinical variables were discarded to guarantee label reliability.
2. **Feature Encoding:**
   - Categorical variables were one-hot encoded (*e.g.*, treatment type, smoking status, genetic mutations).
   - Ordinal or binary attributes were label-encoded (*e.g.*, cancer stage, comorbidities).
   - The output label was mapped to integers 0,1,2.
3. **Feature Scaling:** All continuous variables (*e.g.*, *Age*, *TumorSize*, *FollowUpMonths*, *CancerStage*) were normalized using Z-score scaling (StandardScaler) to stabilize the gradients of the MLP classifier.

The resulting dataset consists of 23 fully numerical features ready for federated training.

FL Configuration

The FL environment was implemented using the Flower framework, simulating a federation of 3 independent nodes (hospital clients). Data were split using Stratified K-Fold partitioning, ensuring that each client retained a representative class distribution while maintaining statistical heterogeneity typical of real clinical settings. Each client reserved 20% of its partition as a local test set.

The simulation was executed using the Flower CLI engine, which automatically orchestrates client/server processes and ensures reproducibility of the full FL pipeline.

Homomorphic Encryption Setup

To guarantee end-to-end confidentiality of model updates, all training rounds were executed with CKKS Homomorphic Encryption, implemented through the TenSEAL library. Server-side parameters were statically defined as follows:

- **Polynomial modulus degree (N):** this parameter determines the size of the polynomial lattice and, consequently, the cryptographic robustness and the ciphertext expansion. A value of 8192 is typically adopted as the standard to ensure a security level of at least 128 bits.
- **Coefficient modulus chain:** $[60, 40, 40, 60]$ $bits$. The values $60bits$ are used for encryption and decryption, while the intermediate values, $40bits$, support arithmetic operations (multiplications) while preserving decimal precision.
- **Global scale:** $2\hat{}$: this scaling factor is required to convert the floating-point values of the neural network weights into polynomial integers that can be efficiently handled by the CKKS scheme.

Clients encrypted their updated model parameters (MLP weights) before transmission, and the central server performed aggregation directly in the encrypted domain using the custom HEFedAvg strategy. Only the aggregated ciphertext was decrypted to obtain the new global model.

Model Architecture

The classification model used by each client was a *Multi-Layer Perceptron (MLP)* with the following structure: *Input of 22–23 features*, *Hidden layers of 64 and 32 neurons* (ReLU activation) and *Output of 3 classes*. The model was trained incrementally using partial_fit, enabling federated updates while preserving optimizer state via warm_start=True.

Table 1
Representation of the results achieved by MLP classifier during our tasks.

Metric	Average Value
Accuracy	0.9240
Precision	0.9215
Recall	0.9240
F1-Score	0.9226

4.2. Evaluation Metrics

In this section, the metrics are introduced. These metrics, which are derived from the Confusion Matrix (CM), play a fundamental role in quantifying the effectiveness of classification models. Accuracy represents the overall proportion of correctly predicted instances in the dataset:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \tag{1}$$

where TP, TN, FP, and FN denote the number of true positives, true negatives, false positives, and false negatives, respectively. Recall measures the fraction of actual positive instances that are correctly identified:

$$Recall = \frac{TP}{TP + FN} \tag{2}$$

Precision quantifies the proportion of correctly predicted positive instances with respect to all predicted positives:

$$Precision = \frac{TP}{TP + FP} \tag{3}$$

The F1 Score computes the harmonic mean of precision and recall, providing a single metric that balances both aspects, particularly useful in scenarios involving class imbalance:

$$F1 = \frac{2}{\frac{1}{\text{Precision}} + \frac{1}{\text{Recall}}} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \tag{4}$$

The F1 score thus consolidates precision and recall into a unified indicator, enabling a more robust assessment of classifier performance in the presence of imbalanced data.

The global evaluation used an aggregated test set combining the 20% local hold-out sets from all clients, following the methodology implemented in the server module.

5. Results

This section presents and discusses the experimental evaluation results of our federated model.

In particular, to assess the performance of the classification model, we adopted Accuracy, Precision, Recall, and F1-Score as evaluation metrics. The results are summarized in Table 1, which reports the performance achieved by the classifier. In particular, the results shown in Table 1, show a high overall Accuracy, confirming that the HEFedAvg strategy succeeds in building a reliable classifier despite the approximation introduced by the CKKS cryptographic scheme. Moreover, the balance between Precision and Recall (highlighted by a weighted F1-Score above 0.92) is particularly important in the medical domain. Indeed a high Precision indicates that the system produces few false positives for a given pathology, reducing the risk of incorrect diagnoses and a high Recall ensures that the system is sensitive, meaning it can detect the majority of actual cases, thereby minimizing false negatives that could lead to missed treatments.

Additionally, confusion matrices are used to illustrate classification behavior during the testing phase. These matrices are shown in Figure 3, providing a visual representation of the model's performance in terms of true positives, true negatives, false positives, and false negatives.

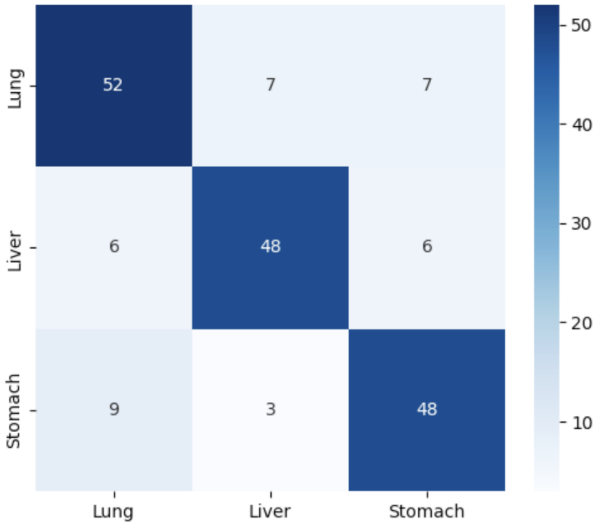


Figure 3: Depiction of confusion matrix values during the evaluation of MLP classifiers. The matrix compares the true labels (Y-axis) with the labels predicted by the model (X-axis) for the three target classes: *Lung*, *Liver*, and *Stomach*.

6. Conclusion

In this work, we presented a privacy-preserving FL architecture that enables collaborative clinical data analysis without compromising patient confidentiality. The proposed framework integrates the Flower federated learning framework with the CKKS homomorphic encryption scheme, ensuring end-to-end protection of model updates during transmission and aggregation. To support encrypted training, we introduced the HEFedAvg strategy, which performs weighted model aggregation directly in the encrypted domain.

The proposed design satisfies several critical requirements for healthcare deployment: (i) *Privacy-by-design:* No raw or intermediate patient data ever leaves the institution; only encrypted model updates are shared, (ii) *Scalability:* Flower-based orchestration supports arbitrary numbers of hospitals or data silos, (iii) *Auditability:* Extensive logging ensures traceability of parameter dimensions, update norms, and training behavior and, (iv) *Hardware compliance:* The framework executes efficiently on CPU-only hospital infrastructures.

Experimental results on a distributed synthetic oncology dataset demonstrated that the proposed solution achieves high diagnostic performance, reaching an accuracy of 92.40% and a weighted F1-score of 92.26%, despite the numerical noise introduced by homomorphic encryption. These results confirm the feasibility of conducting encrypted federated training on tabular clinical data while preserving model convergence and predictive accuracy. Furthermore, the analysis of communication and computational overhead highlighted that, although ciphertext expansion increases network traffic, the overall overhead remains acceptable for realistic clinical deployments.

In summary, our study shows that homomorphically encrypted FL represents a viable and effective approach for secure multi-institution collaboration in the healthcare domain. By combining privacy-by-design principles with robust learning performance, the proposed system provides a practical foundation for future privacy-preserving clinical analytics. It opens the way for secure, large-scale medical data infrastructures.

As future directions, several promising research avenues emerge from this work. First, it will be essential to optimize bandwidth, potentially through gradient compression techniques and more aggressive CKKS packing strategies to reduce transmission overhead. Second, integrating Differential Privacy mechanisms, either during training or at the model-output level, could further protect individual-level privacy and mitigate inference attacks, complementing the confidentiality guarantees provided by homomorphic encryption. Finally, decentralized key management, using threshold homomorphic

encryption or multiparty computation, offers a path toward eliminating the need for a single trusted server, thereby improving robustness and removing central points of failure.

Acknowledgments

This work was supported by the following projects: National Recovery and Resilience Plan, Mission 4 "Education and Research" – Component 2 "From Research to Business" – Investment 1.3, funded by the European Union – NextGenerationEU – CUP: F53C22000740007 – through the project titled "CyberSocial Security," acronym CSS, European Union's Horizon Europe Framework Programme (HORIZON) under Grant Agreement No. 101095119 — A toolkit for dynaMic health Impact analysiS to predicT disability-Related costs in the Aging population based on three case studies of steeL-industry exposed areas in Europe (MISTRAL) and, 'LIFE Project: the itaLian system wide Frailty nEtwork' - local code: T2-AN-12 (CUP D93C22000640001), financed on 'PSC Salute, Traiettoria 2- eHealth, diagnostica avanzata, medical device e mini invasività'.

Declaration on Generative AI

The authors have not employed any Generative AI tools.

References

[1] D. Lofù, P. Sorino, T. Colafiglio, C. Bonfiglio, R. Donghia, G. Giannelli, A. Lombardi, T. Di Noia, E. Di Sciascio, F. Narducci, Morix: Machine learning-aided framework for lethality detection and mortality inference with explainable artificial intelligence in mafld subjects, volume 7, Elsevier, 2025, p. 100176.

[2] Y.-M. Chen, T.-H. Hsiao, C.-H. Lin, Y. C. Fann, Unlocking precision medicine: clinical applications of integrating health records, genetics, and immunology through artificial intelligence, Journal of Biomedical Science 32 (2025) 16.

[3] A. S. Tang, S. R. Woldemariam, S. Miramontes, B. Norgeot, T. T. Oskotsky, M. Sirota, Harnessing ehr data for health research, Nature Medicine 30 (2024) 1847–1855.

[4] P. Sorino, G. Colonna, D. Lofù, T. Colafiglio, A. Lombardi, F. Narducci, T. Di Noia, An explainable machine learning approach for heartbeat classification through signal-based features, in: 2024 IEEE International Conference on Systems, Man, and Cybernetics (SMC), IEEE, 2024, pp. 1–6.

[5] P. Sorino, V. Paparella, D. Lofu, T. Colafiglio, E. Di Sciascio, F. Narducci, R. Sardone, T. Di Noia, A pareto-optimality-based approach for selecting the best machine learning models in mild cognitive impairment prediction, in: 2023 IEEE International Conference on Systems, Man, and Cybernetics (SMC), IEEE, 2023, pp. 3822–3827.

[6] Z. Qu, K. Lin, Z. Li, J. Zhou, Z. Zhou, A unified linear speedup analysis of federated averaging and nesterov fedavg, Journal of Artificial Intelligence Research 78 (2023) 1143–1200.

[7] M. S. Jere, T. Farnan, F. Koushanfar, A taxonomy of attacks on federated learning, IEEE Security & Privacy 19 (2020) 20–28.

[8] J. Shi, W. Wan, S. Hu, J. Lu, L. Y. Zhang, Challenges and approaches for mitigating byzantine attacks in federated learning, in: 2022 IEEE International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom), IEEE, 2022, pp. 139–146.

[9] C. Dwork, Differential privacy, in: International colloquium on automata, languages, and programming, Springer, 2006, pp. 1–12.

[10] B. Knott, S. Venkataraman, A. Hannun, S. Sengupta, M. Ibrahim, L. van der Maaten, Crypten: Secure multi-party computation meets machine learning, Advances in Neural Information Processing Systems 34 (2021) 4961–4973.

[11] A. Acar, H. Aksu, A. S. Uluagac, M. Conti, A survey on homomorphic encryption schemes: Theory and implementation, ACM Computing Surveys (Csur) 51 (2018) 1–35.

[12] M. J. Sheller, G. A. Reina, B. Edwards, J. Martin, S. Bakas, Multi-institutional deep learning modeling without sharing patient data: A feasibility study on brain tumor segmentation, Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries (2018) 92–104. doi:10.1007/978-3-030-11723-8_9.

[13] S. Silva, et al., Federated learning for breast density classification: A real-world experimental study, IEEE Access 10 (2021) 12345–12358. doi:10.1109/ACCESS.2021.3056789.

[14] R. Ramos, et al., Privacy-preserving federated learning for breast cancer prediction, Computers in Biology and Medicine 147 (2022) 105693. doi:10.1016/j.compbiomed.2022.105693.

[15] W. Li, et al., Federated learning for breast histopathology image classification, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, 2020.

[16] A. Roy, et al., Federated learning for brain tumor detection on mri: A multicenter evaluation, IEEE Journal of Biomedical and Health Informatics (2021). doi:10.1109/JBHI.2021.3051234.

[17] M. Rahman, et al., Explainable federated learning for breast cancer diagnosis, Knowledge-Based Systems 251 (2022) 109121. doi:10.1016/j.knosys.2022.109121.

[18] A. Johnson, et al., Federated learning for predicting clinical outcomes from ehrs, Journal of Biomedical Informatics 118 (2021) 103799. doi:10.1016/j.jbi.2021.103799.

[19] N. Rieke, et al., Federated learning for tabular medical data: Challenges and opportunities, Patterns 3 (2022) 100588. doi:10.1016/j.patter.2022.100588.

[20] G. Kaissis, et al., A systematic review of federated learning in healthcare: Challenges and future directions, npj Digital Medicine 5 (2022) 1–19. doi:10.1038/s41746-022-00545-2.

[21] B. Pfitzner, N. Steckhan, B. Arnrich, Federated learning in a medical context: A systematic literature review, IEEE Access 8 (2020) 140699–140719. doi:10.1109/ACCESS.2020.3017316.

[22] T. Li, et al., Privacy-enhancing technologies in federated learning for healthcare: A survey, IEEE Transactions on Big Data (2023). doi:10.1109/TBDATA.2023.3246789.

[23] Q. Yang, et al., A comprehensive survey on privacy-preserving federated learning, ACM Computing Surveys 55 (2023) 1–37. doi:10.1145/3501248.

[24] S. Truex, et al., Security and privacy in federated learning: A survey, IEEE Transactions on Neural Networks and Learning Systems (2022). doi:10.1109/TNNLS.2022.3150421.

[25] J. H. Cheon, A. Kim, Y. Song, Privacy-preserving deep learning via homomorphic encryption: Recent advances and challenges, in: Handbook of Cryptographic Engineering, Springer, 2022, pp. 567–592. doi:10.1007/978-3-030-80737-9_22.

[26] A. Sharma, K. Herath, S. Bagchi, C. Liu, S. Chaterji, Hubs and spokes learning: Efficient and scalable collaborative machine learning, arXiv preprint arXiv:2504.20988 (2025).

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

[28] F. Castro, D. Impedovo, G. Pirlo, An efficient and privacy-preserving federated learning approach based on homomorphic encryption, IEEE Open Journal of the Computer Society (2025).