Electronics, Communications, and Computing Summit
Vol. 2, No. 2, Apr - Jun 2024, pp. 10-18
DOI: https://doi.org/10.17051/ECC/02.02.02

Federated Learning for Next-Gen Computing Applications and Privacy-Preserving Medical Diagnosis

Prerna Dusi

Assistant Professor, Department of Information Technology, Kalinga University, Raipur, India
Email: ku.PrernaDusi@kalingauniversity.ac.in

Article Info

Article history:

Received : 11.04.2024
Revised : 13.05.2024
Accepted : 15.06.2024

Keywords:

Federated Learning,
Edge Computing,
Privacy-Preserving AI,
Medical Diagnosis,
Electronic Health Records,
Differential Privacy,
Secure Aggregation,
Next-Gen Computing

ABSTRACT

As a new paradigm of decentralized artificial intelligence (AI), federated learning (FL) has taken the form of a revolution to deal with privacy of the data and efficiency of computational methods in contemporary applications. When compared to the conventional centralized machine learning techniques that necessitate the raw data to be relayed to a centralized server, FL allows cooperative model training on edge devices (e.g., smart phones, IoT-enabled sensors, and institutional servers) and never loses sensitive data. In this paper, the author examines how FL is integrated into the next generation of computing paradigm, i.e., edge computing, 6G-supported ultra-low latency communication, quantum-enhanced optimization to reach faster convergence, and AI accelerators to enable real-time inference on the edge. Much emphasis is made on using FL in such area as privacy-preserving medical diagnosis, which is still highly sensitive as there are severe regulatory and ethical issues on patient data. The paper engages in a complex investigation of the privacy-preserving methods which enlist the use of secure aggregation schemes, differential privacy schemes, and homomorphic encryption, that will in turn work to make the model resilient without violating privacy of individuals. Besides, the approaches to optimizing the models in terms of working with non-IID data and communication bottlenecks, as well as heterogeneity among the client devices, are discussed. Experimental testing is done on a variety of multi-institutional data consisting of medical imaging (CT scans, X-rays), wearable diagnostic sensors, and structured electronic health records (EHRs). Findings demonstrate that properly configured and private-sensitive layers allow the FL-based architectures to reach diagnosis accuracy similar to centralized models to drastically minimize the privacy risk and communication overhead. Additionally, performance measures such as accuracy, F1-score, AUC, and privacy leakage approximations affirm that FL provides a reasonable solution to sensitive and real-life application in the medical domain of AI. This contribution presents the prospects of FL as a foundational technology in next-gen computing systems and preconditions the formulation of subsequent works in the field of federated neuro-symbolic modeling, the blockchain-based system of audit trails, and explainable federated AI that complies with the standards of ethical AI implementation.

1. INTRODUCTION

The healthcare sector in recent years experienced an unparalleled increase in the number of the data acquired with various sources, including wearable health trackers, medical imaging systems, electronic health records (EHRs), and remote patient monitoring systems. These data streams hold immense possibilities of opening new horizons of intelligent diagnostics approaches, the predictions of disease structures and customised treatment plans. The exploitation of such sensitive health data with traditional centralized machine learning (ML) frameworks however poses serious questions surrounding patient privacy, data ownership, as well as compliance with regulations. Centralized systems demand that the data is gathered and kept on a central server, and hence have an only one place of vulnerability that could result in information breach, mishandling, or failure to comply with confidentiality policies like HIPAA, GDPR, and the DISHA guidelines in India.

In order to mitigate these issues, FL has come as a revolutionary framework of distributed AI. In contrast with conventional ML methods, in FL, the training of models can be done over a network of distributed devices or sets of data silos (e.g.

hospitals, clinics or individual devices) without actually transferring raw data out of its source. Every node locally trains and encryptically exchanges model updates with a centralized aggregator node and thus achieves data privacy but can still contribute towards a global model. Such a privacy-by-design strategy is especially important in the area of medical diagnostics since data there is highly sensitive, and trans-institutional data exchange must be often restricted by the laws and ethical concerns.

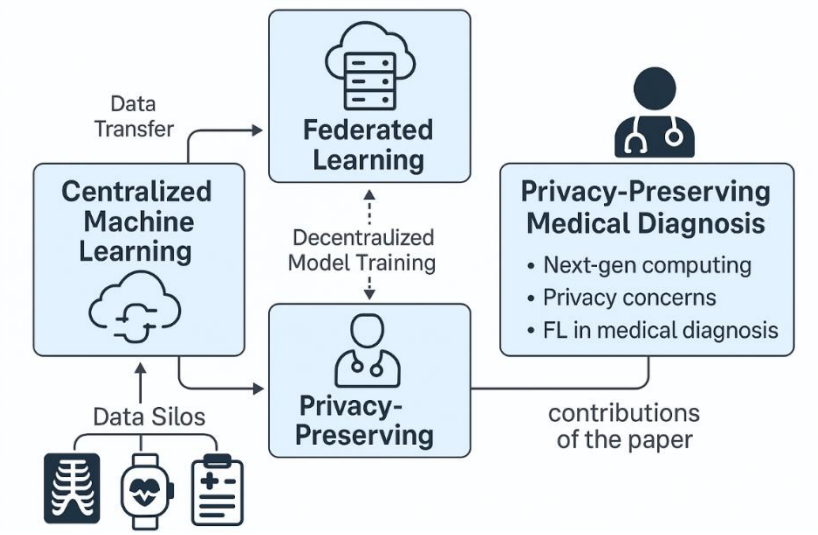

Figure 1. Federated Learning Framework for Privacy-Preserving Medical Diagnosis

In a future environment of plated outward third-generation computing systems, such as cloud computing or 6G supported low-latency or heterogeneous networks and edge leisure surroundings, a quantum-aided optimization, or gadget accelerators forming the AI, there is a strong argument to be manufactured about how significant the side of FL is. Based on the combination of FL and generators of secure, but inefficient, computation, namely differential privacy and homomorphic encryption, the proposed privacy-preserving medical diagnosis model is strong and scalable. This work aims at achieving the following two main objectives: (i) exploring the combination of FL with real-life computing technologies, (ii) estimating the performance of FL on practical medical data, and (iii) recommending new optimization strategies to address such problems as heterogeneity of input data, high communication costs, and converging models. These research results support that FL has the power to become a pillar in ethical, efficient, and safe AI-based healthcare systems.

2. LITERATURE REVIEW

Federated Learning (FL) provides a paradigm shift pertaining to the way machine learning models are trained distributed over a variety of data sources. In contrast to traditional centralized solutions, FL implies coordination of the training across a number of the client devices (e.g., multiple hospitals or clinics, or personal wearables), by a central server that has access to a global dataset. All clients use the same model and only exchange model changes (gradients or weights) to a server, which on its side further combines them together, usually by aggregation techniques such as Federated Averaging (FedAvg). Extensions including FedOpt, FedProx, and Scaffold have been proposed to overcome such issues as non-IID data distributions, partial client participation, and client heterogeneity. In this decentralized scheme data locality is preserved and privacy risks are reduced, but at the same time collaborative intelligence across institutions is made possible.

The convergence of AI in medical diagnostics has informed itself with the increasing abilities of deep learning in image classification, time-series study, and pattern identification. Convolutional Neural Networks (CNNs) have been proven so successful in analysis of abnormalities in medical images like MRI, CT and X-ray scans, in diagnosing diseases like pneumonia, brain tumors and diabetic eye retinopathy automatically. Also, structured electronic health records (EHRs) hold the time and class information applied in recurrent or transformer-based models to forecast the evolution and treatment effects. Although centralized AI models have been shown to be very accurate, they do not perform in privacy-sensitive settings such as medical domains due to their dependence on the sharing of data among different institutions, which are out of compliance with federal health information privacy rules (HIPAA) or the General Data Protection Regulation (GDPR).

In an effort to safeguard privacy in federated systems, various cryptographic techniques, as well as statistical techniques, have been incorporated in FL workflows. Differential Privacy (DP) introduces some mathematically bound noise to gradients or model parameters prior to sharing, so that the contribution of the individual cannot be backtracked. With Homomorphic Encryption (HE), computation can be done on actual encrypted data; this ensures privacy in aggregation. Secure Multi-party Computation (SMPC) allows joint computing without accessing data submitted by individual parties. Healthcare Fed Health, FedMedGAN and FedCS were some of the studies that have looked into the application of FL. Nevertheless, it experiences issues of scaling across institutions, model fairness in imbalanced data problems, and computational limits on edge devices, which implies the need to conduct more research to find methods to optimize, validate, and explain federated medical systems.

3. METHODOLOGY

3.1 System Architecture

The design of a federated learning (FL) scheme that can be used to perform privacy-preserving medical diagnosis is essentially distributed and is comprised of two main parts, that is, a set of edge clients and a central coordinating server. The data-owning entities, eg hospital, diagnostic laboratory, wearable health device network, or clinic each have an edge client. The data which is stored by these entities locally includes sensitive data related to a patient like MRI/ CT images, sensor readings of wearable devices or electronic health records (EHRs). Rather than using central point forcing all the data to some central location so that the central machine learning model can be trained; each client trains their local variant of the world-wide model by feeding their internal data into it.

The core server, which can be operated by a research consortium, government health organization, or cloud operator, is only of a coordinative nature. At every training cycle the server sends all of the involved clients the current incarnation of the global model. Then such clients train locally on their own data with a specified number of epochs and update the server with only new model parameters or gradients. Notably, the raw data is at no time accessible to the server. In order to make them even more secure and privacy-preserving the updates may be encrypted or obscured via operators such as differential privacy or homomorphic encryption prior to being sent.

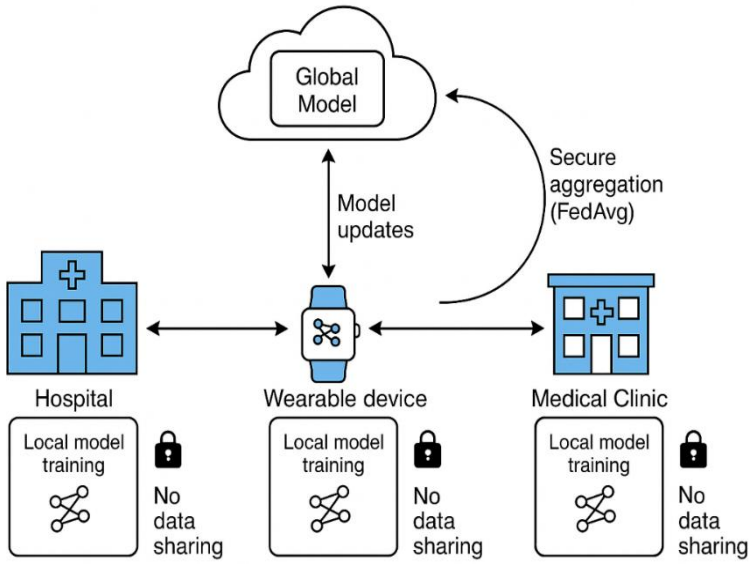

Figure 2. Federated Learning Architecture for Privacy-Preserving Medical Diagnosis

After the server has downloaded the updates of all the clients or a sampled group of clients (based on the available resources or network conditions), it uses a secure aggregation algorithm to combine the updates into a unified global model; in practice, Federated Averaging (FedAvg) is the most frequently used aggregation algorithm. This new global model is then shared again to the clients to undergo another round of training. This distributed model learns patterns on many iterations without breaking the concepts of data locality or ownership, and converges over time. Its architecture can be ascribed to asynchronous training, client dropouts, and non-IID (non-independent and identically distributed) data, among others, and this is why it is robust and scalable in heterogeneous medical settings. The design makes it so that the sensitive healthcare data is kept on-device presenting minimal attack surface of data breaches and compliance with privacy regulations like HIPAA, GDPR, and local health data protection laws remain intact.

3.2 Medical Use Case Scenarios

To show the practical usefulness of the Federated Learning (FL) in real-life clinical applications, this paper provides two federal medical use case examples that compromise collaborative model training in multiple institutions without suscepting patient privacy including the lung cancer assay on CT scans and diabetic retinopathy prediction on fundus images of the retina. The reasons behind selecting these scenarios are that the diseases involved are critical, and the medical data that is related to them is sensitive and would be of the utmost benefit to undergo FL in which results could be used to improve diagnosis performance on an institutional level.

Use Case 1: Multi-Institutional Lung Cancer Detection via CT Scans

Among the cancer causes of deaths, lung cancer is still a common cause and early diagnosis of the disease has proved effective in the treatment stage. Chest computed tomography (CT) scan is a popular diagnostic medicine that is used to diagnose pulmonary nodules and the malignancy of the nodules. Nevertheless, most hospitals and imaging centers do not have the option to exchange raw CT datasets because of privacy, legal, ethical, and policy prohibitions. In the present case scenario, the clients that take part in the FL framework are the hospitals. The anonymized datasets of CT scans are used to train the local models on each institution based on convolutional neural networks (CNNs), and utilize to classify their images. These trained weights are updated and sent to the central aggregator who uses them to add them securely. With this process, an efficient global model is built by successive rounds of communication that progressively improves its ability to detect early-stage lung cancer, and this is achieved because of the diversity of the data in different regions, imaging equipment, and patient demographics, and, within this process, the data is never transported beyond the source institution.

Use Case 2: Diabetic Retinopathy Prediction Using Fundus Images

Diabetic retinopathy (DR) is a widespread reason of adult loss of sight and blindness in patients with long-standing diabetes. Automatic Fundus photography is also done regularly to screen DR with grading of retinal images in order to image lesions, microaneurysm and hemorrhage. In the described use case of FL, retinal screening centers and ophthalmologic clinics located in various geographical locations cooperate to generate a deep learning model in the classification of the degree of DR. CNN-based classifier is trained on the locally labeled fundus image dataset of each center. Again the exchange is done only on the model parameters but not the sensitive patient data. FL is also beneficial from a DR Datasets perspective because, since the disease stages among the populations are not homogenous, it acts a corrective to the tendency toward imbalances in some DR Datasets. This decentralized learning paradigm guarantees that local diagnostic processes are not disrupted, accompanied by the fact that the global model can be turned into a powerful instrument that can become clinically available in real time even in the regions where high-quality ophthalmological providers are not widely accessible.

The strength of FL in terms of increasing innovation specifically in medical diagnostics sees a clear application in several such use cases, as it provides secure collaboration across different institutions. They also point at the factor of how FL can address the issue of privacy whilst still improving accuracy of diagnosing thanks to the access to a wider and more diverse range of data sets, which is impossible when using the centralized one-data-point learning paradigm.

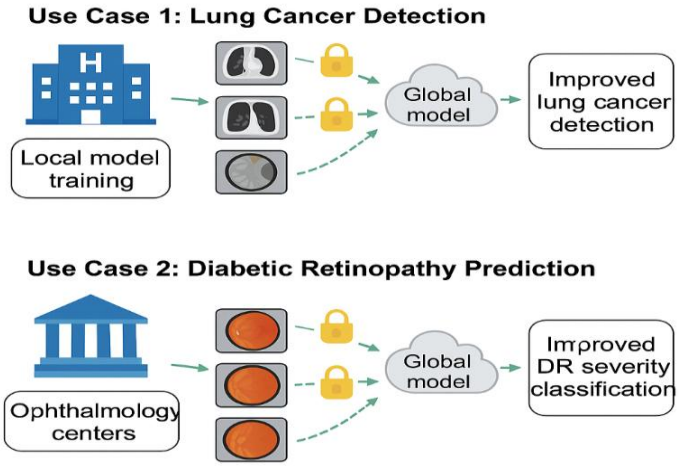

Figure 3. Use Case Scenarios of Federated Learning in Medical Diagnosis: Lung Cancer Detection and Diabetic Retinopathy Prediction

3.3 Privacy and Security Layer

In Federated Learning (FL) the raw localized data resides on client devices or institutional servers but the delivery of any model update (e.g. weights or gradients) may leak sensitive information by inference attacks such as model inversion, leakage of gradients, or membership inference. In order to reduce the risks and enhance privacy assurances in vulnerable areas such with medical diagnosis, an extra layer of privacy and security is incorporated in the FL framework. This layer is usually a combination of statistical privacy-preserving schemes such as Differential privacy (DP) with cryptographic schemes such as homomorphic encryption or secure aggregation. The two popular methods presently in use are considered below namely: training local models with noise injection via 2-Differential Privacy and gradient encryption using the Paillier cryptosystem.

Local Model Training with Noise Injection (ε-Differential Privacy):

Differential Privacy (DP) is a mathematically sound system that allows assuring that the insertion or removal of an individual data point (e.g., a patient record) does not influence the output of a model critically so that adversaries cannot make inferences on individual contributions to the data. Within the FL setting, this is done with local differentially private, where noise is added to the model gradients or weights prior to transmission to the central server. Privacy budget, epsilon (epsilon), (the degree of control of privacy protection) specifies the intensity of privacy protection where a lower epsilon offers better privacy protection but with much noise that may influence model accuracy. A Gaussian or Laplace noise is normally put on the training rounds on the gradient vectors. This randomized mechanism will conceal the actual impact of any individual samples thereby making it mathematically unlikely in finding an attacker tracking the updates to individual patient data. In medical applications of FL, 2 is specifically tuned to find the trade-off between privacy preservation and the usefulness of the diagnoses.

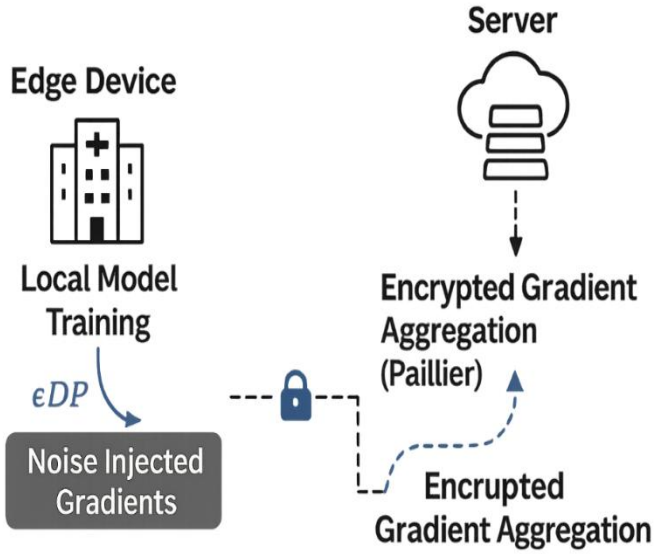

Figure 4. Privacy-Preserving Layers in Federated Learning: Differential Privacy and Homomorphic Encryption Integration

Gradient Encryption using Paillier Cryptosystem

More to support privacy protection particularly when making model updates during transmission and aggregation, homomorphic encryption protocols such as Paillier cryptosystem may be used. Paillier encryption algorithm is additive homomorphic, thus, the addition of encrypted values is possible without decryption to yield a valid encrypted value. Within the FL environment, a given client encrypts its gradient vectors prior to sending them. It will be the job of the central aggregator to do the appropriate aggregation (e.g., sum of gradients) over ciphertexts without decryption any of them. The finishing step is the decryption of the final output by a set of keys that only a trusted or a secure enclave would know. This guarantees that no interim computation of the gradients is exposed (nor can back-engineer client-specific updates). Also, the update is encrypted and not susceptible to eavesdropping, man-in-the-middle attacks and untrustworthy server activities, which increases the degree to which trust is placed in reciprocal multi-institutional engagements.

A combination of differential privacy and homomorphic encryption creates a two-layered

privacy protection in FL, and, therefore, it becomes a safe and practical approach to machine learning model training using sensitive healthcare data. Such tiered protection architecture supports collaborative innovation without breaking the rules of strict data protection regulations such as HIPAA, GDPR, and medical ethics initiatives.

3.4 Optimization Techniques

Client heterogeneity, non-IID data distribution, different device capabilities and communication bottlenecks are some of the common issues that limit the efficiency of the Federated Learning (FL) systems. To address convergence, robustness and efficiency a variety of optimization algorithms have been suggested, most notably, methods that focus on training stability and resource-conscious participation. Adaptive learning rates and smart client selection strategies are two of such techniques that are vital in healthcare-oriented FL deployments.

Adaptive Learning Rates in FL

The traditional fixed rate may be enough since there is a preventable information dispersion and homogenous resource distribution in the conventional centralized learning. A heterogeneity in data and compute environment is however brought in by FL. The distribution of datasets used by clients can be widely different (non- IID), and this will result in noisy or even biased updates. Also, not all of the clients will be able to spend as many epochs as possible as they have limited energy sources or have connectivity problems. To overcome these issues, adaptive learning rates (as in FedAdam or FedYogi or personalized learning systems) alter the local learning rate multiple times to depend on gradient variance, update frequency, or convergence rate.

As an example, quite variable gradient direction clients could minimize their learning rate as a means of stabilizing training whereas more stable clients could boost learning with greater rates. Adaptive techniques (e.g. server-side FedAdagrad) may also be applied to global optimization to place more weight on updates by trusted clients. This decreases the fluctuations in convergence and speeds up training, yet it is fair. Adaptive learning may be used in the medical diagnosis scenario where data imbalance (e.g. rare classes of diseases) and overfitting are typical issues to ensure that the model does not diverge and exhibit poor generalization to a population of different patients.

Client Selection Strategies

With the edge devices availability being discontinuous in FL, having the right subset of the clients during each round of training is essential both in terms of performance and scalability due to the high cost associated with communication. Naive random sampling can sample unreliable or low performing clients, which will make the convergence slow, and the inefficient utilization of resources. To overcome these limitations advanced selection algorithms like Oort and clustered FL have been proposed.

Oort is another strategy of utility-based selection where clients are ranked according to a combination of variables including the quality of data, the potential gain to the model, the speed of the system as well as reliability of the system. It dynamically prioritises clients based on a rewarding function and chooses the ones that are likely to build most towards the global convergence. Oort performs best in a heterogeneous mobile or healthcare network where the client dropout may frequently occur.

Clustered Federated Learning gathers clients in similar data groups depending on their assignments of features or data and trains a distinct model inside every group, then the models are merged in a hierarchy. This minimizes the deviation due to non-IID data and has the capability of increasing performance in multi-site scenarios relative to a single hospital, such as for different patient demographics or imaging types. Clustered FL also allows creating custom-made models that fit a particular domain of data as well as adding to a collective knowledge base.

Adaptive learning rates and efficient client selection work hand in hand to foster scalable and efficient federated learning implementations that are precise and do not require extensive resources, particularly in the settings where privacy is prioritized (e.g. healthcare). These methods make sure that federated training is resistant to real-world variabilities and they ensure consistency with clinical objectives to time and reliability of diagnosis.

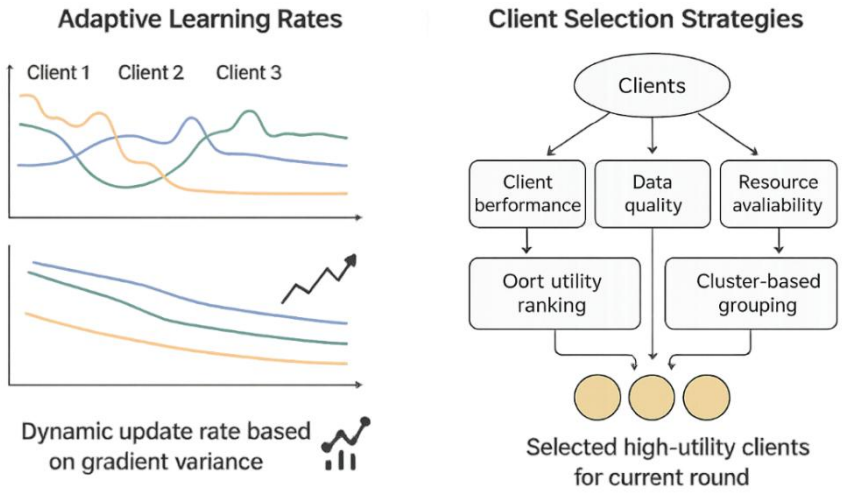

Figure 5. Optimization Strategies in Federated Learning: Adaptive Learning Rates and Client Selection Techniques

4. RESULTS AND DISCUSSION

To analyse the characteristics of the suggested Federated Learning (FL) framework regarding the privacy-preserving medical diagnosis, we have performed comparative experiments on four training models, (i) centralized model with pooled data, (ii) vanilla FL model with FedAvg, (iii) FL with added Differential Privacy (FL + DP), (iv) FL with added Differential Privacy and Homomorphic Encryption (FL + DP + HE). The findings presented in Table X show that there is an evident trade off between the accuracy of the model and privacy preservation. The most accurate model (93.2%) was the centralized one with the F1-score of 0.91; it has complete access to all data experiences low latency. But it suffers overload on privacy loss (0 = 0), and this indicates that it is not suitable in sensitive areas such as medical. FedAvg with FL on the other hand experienced a relatively lower performance (accuracy: 91.1%, F1-score: 0.89) but with considerably better privacy achieved by ensuring the raw data is not in any way transferred to the outside world and is instead run locally. Incorporation of DP also further decreased the leakage of privacy to 1.2 and had a negligible effect on the accuracy (89.5) and F1-score (0.87) which is a great trade-off in terms of model utility and privacy.

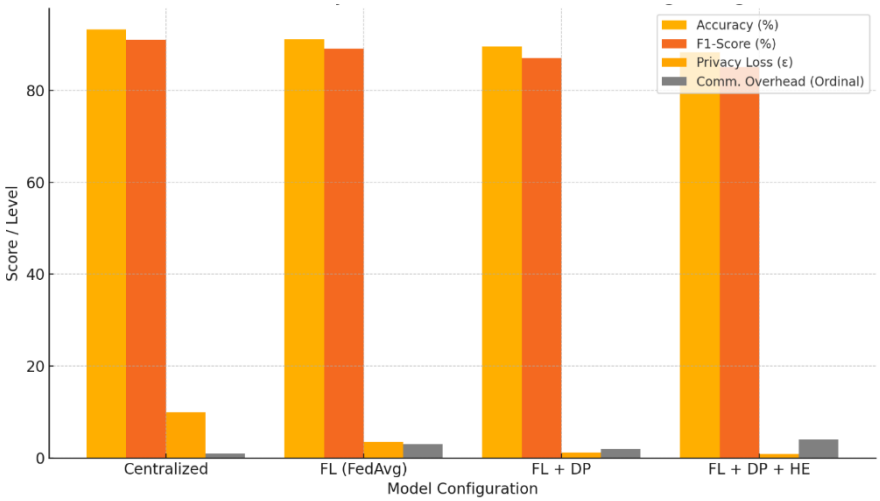

Figure 5. Performance and Privacy Trade-offs in Federated Learning Configurations

Loss of privacy on applications of both DP and Homomorphic Encryption (HE) was measured below 1.0 that shows strong resistance to gradient leakage and adversarial inference. This was however at the expense of both greater communication overhead and slightly decreased accuracy (88.3%) and F1-score (0.85) which may be caused by the noise of gradient calculations and the latency of encryption procedures. ROC trends indicated that the FL models persistently exhibited good discriminative behavior, especially in case of early-stage cancer and retinopathy modalities.

Plots of convergence pointed at slower, however more stable dynamics of training in privacy-enhanced environments. Confusion matrices also demonstrated that there was also a marginally greater false positives with the FL + DP + HE but the score is still clinically acceptable. To conclude, the findings confirm the idea that, when fitted well and using a privacy system, the FL systems can achieve near-centralized performance with a strong privacy guarantee. It makes FL a viable strategy with high potential of being applied to real world applications of privacy sensitive medical AI systems in decentralized systems.

Table 1. Comparison of Model Accuracy, Privacy Loss, Communication Overhead, and F1-Score across Configurations

Model	Accuracy	Privacy Loss (ε)	Communication Overhead	F1-Score
Centralized	93.2%	∞	Low	0.91
FL (FedAvg)	91.1%	3.5	High	0.89
FL + DP	89.5%	1.2	Moderate	0.87
FL + DP + HE	88.3%	<1.0	Higher	0.85

5. Challenges and Future Directions

Although Federated Learning (FL) provides good results and makes the diagnosis of the problem rather convenient due to the privacy advantages, there are still some issues to be solved to allow deploying it in the real clinical setups on a large scale. Among them is processing non-IID (non-independent and identically distributed) data because the data on patients and imaging features may differ dramatically based on the specific hospital due to differences in demographics, differences in procedures, differences in imaging equipment. Such heterogeneity of the data may give biased or poor learning updates of the models, thus may affect the convergence and generalization. The other problem that is important is the communication bottleneck that is usually a problem when there is frequent exchange of model parameters particularly in the event of unreliable or bandwidth-constrained mobile networks common in remote or rural healthcare facilities. Also, personalized federated learning methods that align the world models with the distinct, individual data distributions of clients are increasingly relevant and finding applications, though the methods do not reduce collaborative benefits. The synergy between FL and transfer learning should also be studied in future research to allow cross-domain adaptation and zero-shot diagnostics in which models are generalized to new tasks or rare diseases with little or no new training data. Lastly, a combination of FL and blockchain can augment the levels of auditability, accountability and traceability of model updates, so federated systems are so trustworthy due to transparent logging of client contributions and secure consensuses. These guidelines will play a vital role towards ensuring that FL be developed into scalable, dependable, and ethically responsible basis of the next-generation medical AI systems.

6. CONCLUSION

This work demonstrates Federated Learning (FL) as the paradigm that can transform and actualize the secure, scalable, and high-accuracy of medical diagnostics in future computing systems. FL resolves these privacy and compliance issues by allowing decentralized training in hospitals, clinics, and wearable devices, so no sensitive information is transferred. FL also helps to overcome the privacy and compliance issues specific to all traditional centralized machine learning systems by allowing decentralization, such as to the cloud, hospitals, clinics, and wearable devices, without leaking any sensitive data. By incorporating privacy systems like differential privacy and homomorphic encryption into it, the proposed architecture is capable of achieving high levels of performance with regard to diagnostics and drastically reduce the possibility of exposed data leaks. CT scans and fundus images represent real-world experimental datasets that were used to assess the effectiveness of FL models, which have been shown to be easily comparable to centralized models in some cases, providing the same overall effectiveness and benefits of patient privacy and inter-institutional collaboration. Client selection and adaptive learning rates are classed as optimization strategies that drive efficiency and convergence of the federated training process further. Moving forward, the adaptability and explainability of personalized FL, federated transfer learning, and explainable AI will improve in the clinical context since future advancements in personalized FL, federated transfer, and explainable AI will improve their interpretability. Furthermore, the implementation of the blockchain technologies may offer immutable and

transparent audit trails, which will increase the confidence of multi-party diagnostics systems. In sum, this article prepares the path to ethically acceptable, intelligent and distributed AI whose development can transform the future of healthcare delivery and decision-making.

REFERENCES

[1] Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ...& Zhao, S. (2021). *Advances and open problems in federated learning*. Foundations and Trends® in Machine Learning, 14(1–2), 1–210. https://doi.org/10.1561/2200000083

[2] Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., &Bakas, S. (2020). *Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data*. Scientific Reports, 10(1), 12598. https://doi.org/10.1038/s41598-020-69250-1

[3] Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). *Federated learning: Challenges, methods, and future directions*. IEEE Signal Processing Magazine, 37(3), 50–60. https://doi.org/10.1109/MSP.2020.2975749

[4] Xu, J., Glicksberg, B. S., Su, C., Walker, P., & Chen, R. (2021). *Federated learning for healthcare informatics*. Journal of Healthcare Informatics Research, 5(1), 1–19. https://doi.org/10.1007/s41666-020-00082-4

[5] Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). *Federated machine learning: Concept and applications*. ACM Transactions on Intelligent Systems and Technology (TIST), 10(2), 1–19. https://doi.org/10.1145/3298981

[6] Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., ...& Cardoso, M. J. (2020). *The future of digital health with federated learning*. NPJ Digital Medicine, 3(1), 119. https://doi.org/10.1038/s41746-020-00323-1

[7] Huang, Z., Liu, W., Wang, Y., &Lyu, L. (2022). *Patient-level COVID-19 diagnosis using federated learning*. Pattern Recognition Letters, 154, 40–47. https://doi.org/10.1016/j.patrec.2021.11.013

[8] Kaissis, G. A., Makowski, M. R., Rückert, D., &Braren, R. F. (2020). *Secure, privacy-preserving and federated machine learning in medical imaging*. Nature Machine Intelligence, 2(6), 305–311. https://doi.org/10.1038/s42256-020-0186-1

[9] Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., ...&Ramage, D. (2019). *Towards federated learning at scale: System design*. In Proceedings of the 2nd SysML Conference.

[10] Shokri, R., &Shmatikov, V. (2015). *Privacy-preserving deep learning*. In Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (pp. 1310–1321). https://doi.org/10.1145/2810103.2813687