American Journal of
Computing and Engineering
(AJCE)

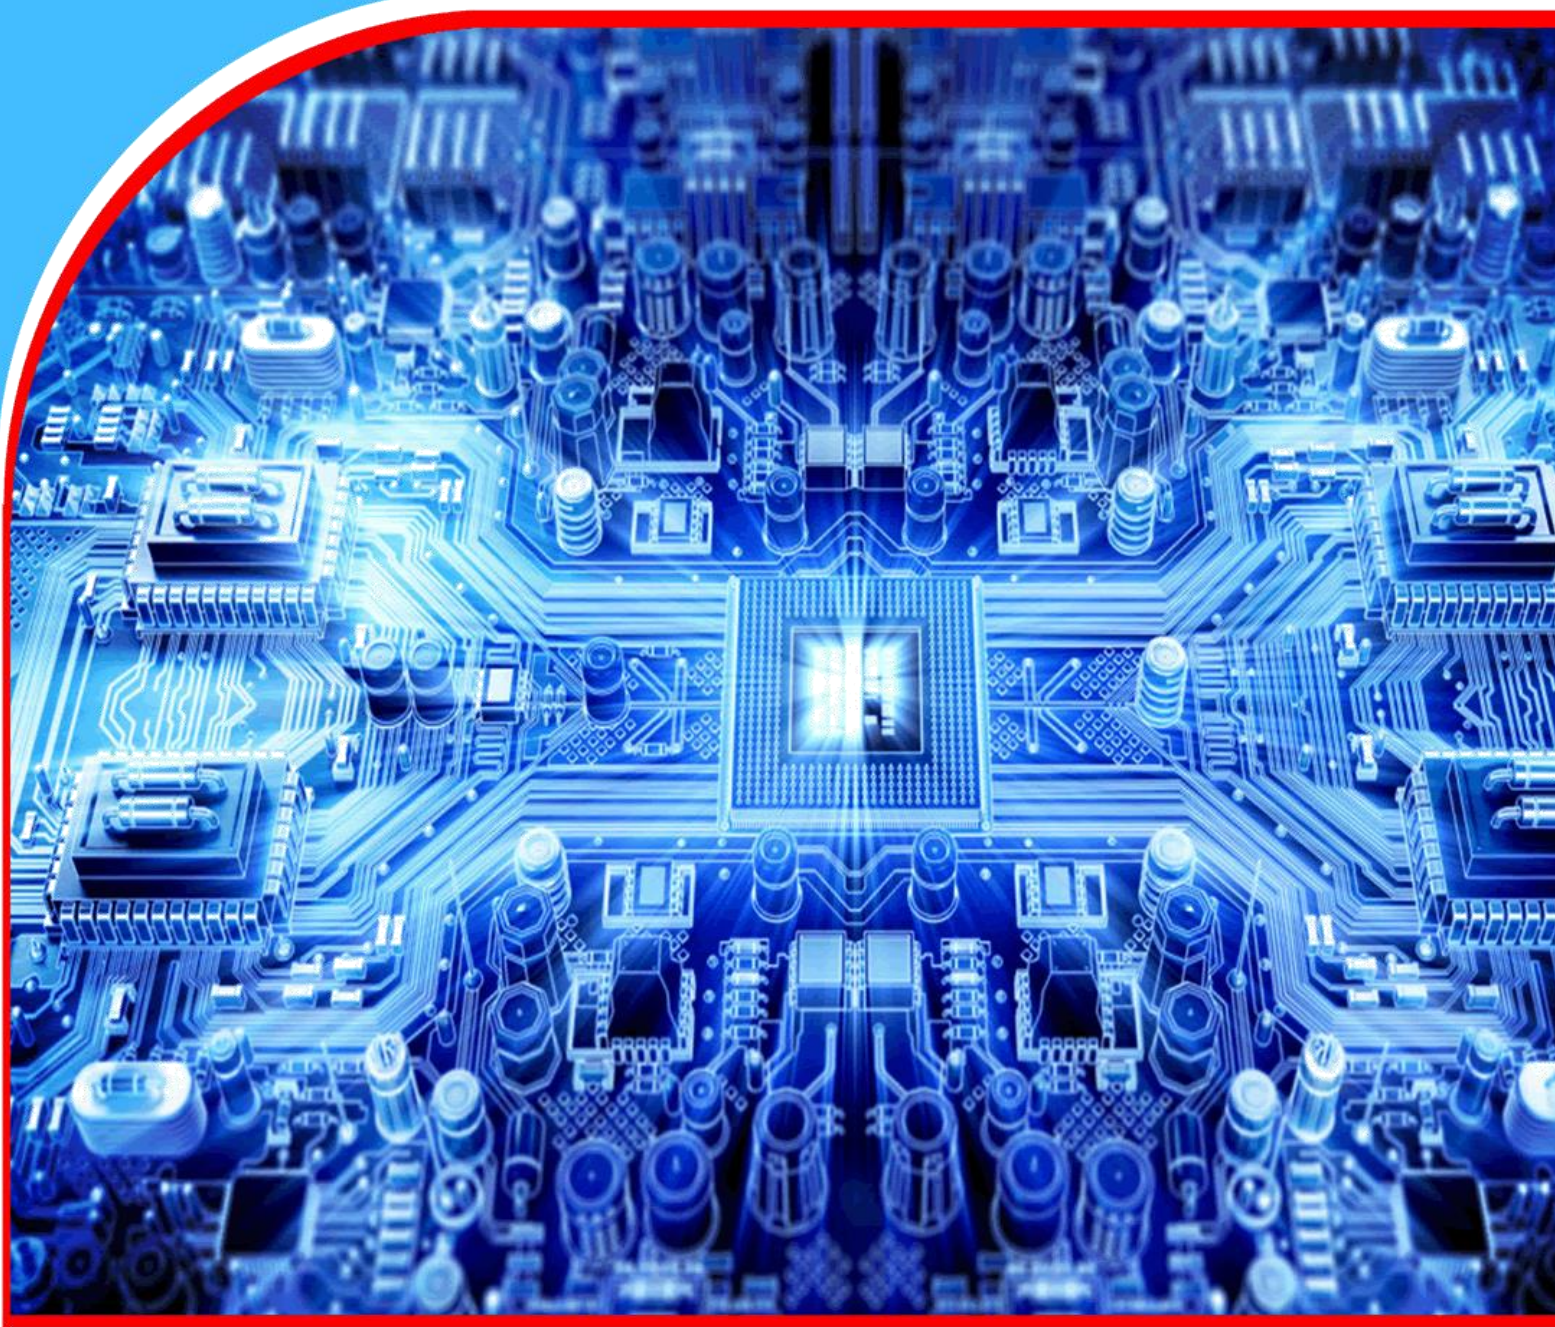

Federated Learning for Healthcare: Balancing Data Privacy and Model Accuracy

Nishchai Jayanna Manjula, Kiran Randhi, Srinivas Reddy Bandarapu

American Journal of Computing and Engineering
ISSN 2790-5586 (Online)
Vol.6, Issue 1, pp 111 – 122, 2023

www.ajpojournals.org

Federated Learning for Healthcare: Balancing Data Privacy and Model Accuracy

Nishchai Jayanna Manjula[1*], Kiran Randhi[2], Srinivas Reddy Bandarapu[3]

[1]Senior Solutions Architect, Amazon Web Services

[2]Principal Solutions Architect, Amazon Web Services

[3]Principal cloud Architect DigiTech Labs

Article history

Submitted 13.09.2023 Revised Version Received 16.11.202 Accepted 20.12.2023

Abstract

Purpose: Federated Learning (FL) is transforming the way machine learning models are trained by allowing institutions to collaborate without sharing sensitive data. This is especially valuable in healthcare, where patient records are often stored separately across hospitals and research centers. This decentralized approach allows healthcare providers, researchers, and organizations to leverage collective intelligence from distributed datasets, leading to advancements in diagnostics, treatment personalization, and patient outcomes.

Materials and Methods: However, the adoption of FL in healthcare is not without challenges, particularly in balancing the dual objectives of preserving data privacy and maintaining model accuracy. In this article, we explore how FL is being applied in healthcare, examining the balance between protecting patient privacy and ensuring high model accuracy. We review recent advancements in FL, focusing on privacy-preserving techniques such as differential privacy, secure multi-party computation, and homomorphic encryption, and their impact on model performance.

Findings: Through a comprehensive analysis of case studies and empirical research, we highlight the potential of FL to revolutionize healthcare applications, including medical imaging, electronic health records (EHR) analysis, and genomic research. We discuss recent advancements, key challenges, and innovative solutions, drawing insights from various studies.

Implications to Theory, Practice and Policy: Finally, we highlight future directions and provide practical recommendations for researchers and professionals looking to implement FL in medical settings.

Keywords: *Federated Learning, Machine Learning, Healthcare*

American Journal of Computing and Engineering
ISSN 2790-5586 (Online)
Vol.6, Issue 1, pp 111 – 122, 2023

INTRODUCTION

The healthcare industry has seen a major shift with the rise of Artificial Intelligence (AI) and Machine Learning (ML). These technologies have enhanced everything from diagnostics to treatment plans, helping doctors make quicker and more precise decisions. AI-driven tools are detecting diseases earlier, analyzing medical images with greater accuracy, and even predicting patient outcomes.

However, to train these advanced models, a large amount of data is needed. The problem is that patient information is highly sensitive and protected by strict privacy laws such as HIPAA (Health Insurance Portability and Accountability Act) [1] and GDPR (General Data Protection Regulation) [2]. This is necessary for data safety as healthcare industry has been of central point for several cyber attacks in past few years causing millions of dollars of loss and breaching of private data [3]. These regulations mentioned earlier, while essential for safeguarding privacy, also create barriers to data sharing. Hospitals and research centers often work in isolation, making it difficult to build AI models that benefit from a diverse range of medical cases.

This is where Federated Learning (FL) steps in as it makes possible for multiple institutions to collaborate on AI model training without sharing raw patient data. FL allows models to learn from decentralized data sources and sharing of only the model updates ensuring privacy is maintained. This approach opens the door for safer, more effective AI-driven healthcare solutions while addressing the critical need for data security.

Need for Privacy of Healthcare Data

Healthcare is one of the most data-driven industries, producing an enormous amount of sensitive information every day. From patient records and diagnostic reports to treatment histories and genetic data, this information plays a crucial role in delivering personalized care, advancing medical research, and improving public health. However, because of its highly sensitive nature, healthcare data is also a major target for cyberattacks, which have become more frequent and sophisticated in recent years. According to a report by IBM, the average cost of a data breach in the healthcare industry reached $10.1 million in 2023, the highest among all sectors [4]. Anthem, one of the largest health insurance companies in the U.S., suffered a massive data breach that exposed the personal information of nearly 78.8 million individuals. The stolen data included names, Social Security numbers, and medical IDs, making it one of the largest healthcare data breaches in history [5]. UHS, a leading healthcare provider in the U.S., experienced a ransomware attack that disrupted operations across its 400 facilities. The attack forced the organization to revert to manual processes, delaying patient care and causing significant financial losses. Thus, protecting this data is not just about maintaining privacy it's essential for ensuring trust in the healthcare system and safeguarding patient well-being and to address these challenges, healthcare organizations have adopted various privacy-preserving techniques, including encryption, access controls, and anonymization. But these traditional privacy-preserving methods have their own limitations, especially when it comes to data sharing and collaboration in healthcare. Encryption and access controls don't fully mitigate the risks of sharing data across institutions, and anonymization can weaken the usefulness of data for research. The rise of AI and machine learning has made this issue even more complex, as these models require large, diverse datasets to be effective. However, sharing such data increases security risks and regulatory challenges, highlighting the urgent need for privacy-preserving solutions that support secure collaboration.

American Journal of Computing and Engineering
ISSN 2790-5586 (Online)
Vol.6, Issue 1, pp 111 – 122, 2023
www.ajpojournals.org

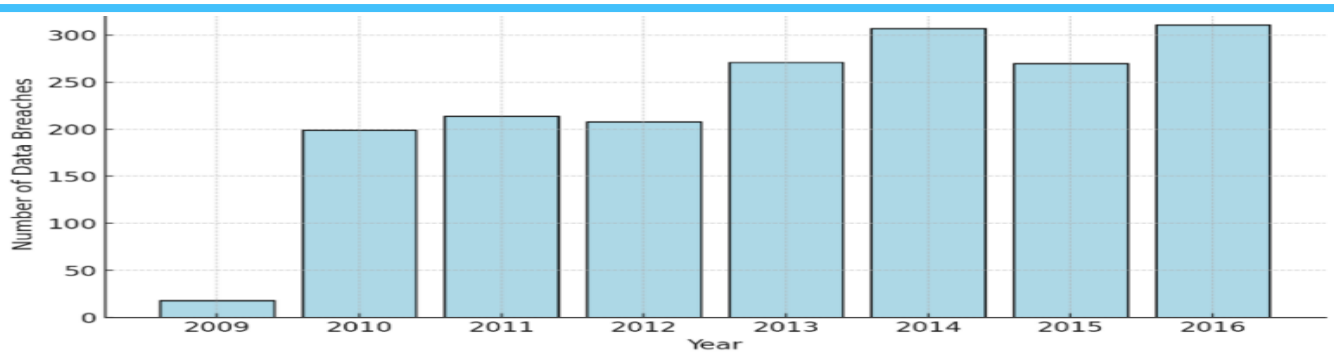

Figure 1: Healthcare Data Breaches: Annual Trend from 2009 to 2016

Federated Learning: A Paradigm Shift for Healthcare Industry

Federated Learning (FL) allows multiple institutions, such as hospitals and clinics, to train a machine learning model collaboratively without sharing sensitive patient data. Instead of sending raw data to a central server, each institution trains the model locally and only shares updates that improve the model. A central coordinating server (often managed by a trusted authority or AI provider) creates an initial machine learning model. This model is untrained or pre-trained on publicly available datasets to establish a starting point. The central server then distributes this initial model to participating hospitals, research institutions, or healthcare providers. Each participant receives a local copy of the model but does not share any of its patient data. Each institution trains the model on its own private dataset. For example, a hospital with lung cancer patient records will train the model using its medical imaging data, while another hospital specializing in diabetes will train the same model on blood glucose patterns.

The training process involves:

- Feeding local patient data into the model.
- Adjusting the model's internal parameters (weights) based on patterns in the data.
- Improving the model's ability to recognize specific healthcare conditions.

Since all training happens locally, patient data never leaves the institution, preserving privacy.Once training is complete, each institution extracts only the updated model parameters (such as gradients or weight adjustments) rather than the actual patient data.The central server collects updates from multiple institutions and aggregates them into a single, improved global model. This process is known as Federated Averaging [8].

This way, Federated Learning allows healthcare institutions to collaborate on AI model training without sharing patient data, improving accuracy while preserving privacy. It enhances disease detection and personalized treatment by learning from diverse datasets across hospitals, ensuring AI models are more adaptable and unbiased. Additionally, FL aligns with data protection laws like HIPAA and GDPR, reduces cybersecurity risks by keeping data decentralized, and empowers smaller medical institutions to contribute to AI advancements without exposing sensitive records, making healthcare AI more inclusive and effective.

Objective and Scope

This article aims to explore the application of FL in healthcare, with a particular focus on balancing data privacy and model accuracy. We will review recent literature, identify key challenges, and discuss potential solutions and try to answer the the current state of Federated Learning in healthcare, and how does it address the challenges of data silos and privacy regulations and the key trade-offs between data privacy and model accuracy in Federated Learning, and how do they impact healthcare applications.

Also, the privacy-preserving techniques are currently used in Federated Learning, and how effective are they in healthcare settings.

American Journal of Computing and Engineering
ISSN 2790-5586 (Online)
Vol.6, Issue 1, pp 111 – 122, 2023
www.ajpojournals.org

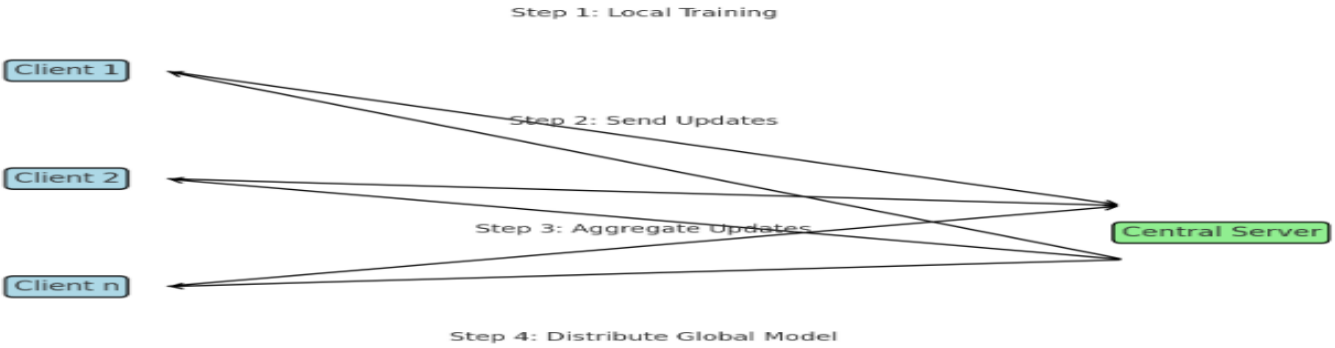

Fig. 2 Schematic Diagram of Federated Learning

THE PRIVACY-ACCURACY TRADE-OFF IN FEDERATED LEARNING

By tackling one of the most important issues of our day data privacy, Federal Learning (FL) has reshaped the fields of artificial intelligence and machine learning. FL achieves a fine balance between protecting user privacy and attaining high model accuracy by permitting several parties to work together on model training without exchanging raw data. This balance is not without difficulties, though, as researchers and practitioners have to carefully assess the trade-off between model performance and privacy preservation. At its core, FL ensures that sensitive data remains decentralized, residing locally on the devices or servers where it was originally generated. Each participating node trains a local model using its data and shares only model updates, such as gradients, with a central aggregator. This architecture inherently reduces the risk of exposing raw data, thus enhancing privacy. However, the use of model updates instead of raw data is not entirely foolproof. While FL offers significant privacy benefits, it also introduces challenges related to model accuracy. The decentralized nature of FL can also lead to issues such as:

Non-IID Data: Healthcare data is often non-independent and identically distributed (non-IID) across institutions, which can degrade model performance [9].

Communication Overhead: Frequent communication between local devices and the central server can lead to latency and bandwidth issues [9].

Model Heterogeneity: Different institutions may use different data formats, feature sets, or model architectures, complicating the aggregation process [10].

Common Privacy-Preserving Techniques in Federated Learning

FL is not completely resistant to attackers, even with its inherent privacy-preserving structure. Reconstruction attacks and membership inference attacks are vulnerabilities that allow malicious actors to take advantage of model modifications to infer private information about the underlying data. Addressing these risks is a necessity in industries where data sensitivity is critical, such as healthcare, banking, and education. Therefore, it is essential to integrate advanced privacy-preserving methods into FL in order to protect user data and uphold system credibility. Various techniques have been developed to enhance the privacy of FL, ensuring that even shared model updates reveal minimal information about individual datasets. Here are some widely adopted approaches:

Differential Privacy introduces statistical noise to model updates before they are sent to the central server. By adding noise, it becomes exceedingly difficult to deduce specific data points from the updates, even if an adversary gains access to them. DP provides strong theoretical guarantees of privacy but can also affect the accuracy of the model if excessive noise is added [11].

SMPC enables multiple parties to collaboratively compute a function over their inputs while keeping these inputs private. In the context of FL, SMPC ensures that model updates are encrypted and only the

aggregated result is decrypted by the server. This approach provides robust privacy but comes with computational and communication overhead [12].

Homomorphic Encryption allows computations to be performed directly on encrypted data without requiring decryption. In FL, clients can encrypt their model updates, send them to the server, and the server can perform aggregation on the encrypted values. This ensures that raw updates remain inaccessible. However, HE can significantly increase computational requirements.

Impact of Privacy-Preserving Techniques on Model Accuracy

While privacy-preserving techniques are essential for mitigating risks, they introduce trade-offs that impact model accuracy. Techniques like Differential Privacy inherently reduce the signal-to-noise ratio in the model updates, potentially degrading the model's ability to generalize. The extent of accuracy loss depends on the level of noise added; higher privacy guarantees often come at the cost of reduced accuracy. Similarly, SMPC and HE significantly increase the computational burden and communication latency, which can indirectly affect model accuracy. For instance, longer training times due to computational delays may result in incomplete convergence within practical timeframes. Techniques that protect privacy may result in aggregated updates that are less helpful; for instance, noised or encrypted gradients may mask important patterns that the central server could use to optimize the global model.

American Journal of Computing and Engineering
ISSN 2790-5586 (Online)
Vol.6, Issue 1, pp 111 – 122, 2023
www.ajpojournals.org

TABLE 1: COMMON PRIVACY PRESERVING TECHNIQUES IN FL

Technique	Description	Advantages	Limitations
Differential Privacy (DP)	Adds controlled noise to model updates or outputs to prevent inference attacks.	Strong theoretical privacy guarantees; prevents re-identification of data.	Noise can degrade model accuracy; requires careful tuning of privacy parameters.
Secure Multi-Party Computation (SMPC)	Enables multiple parties to compute functions over their data without sharing it.	No raw data is exposed; secure aggregation of model updates.	High computational and communication overhead; scalability challenges.
Homomorphic Encryption (HE)	Allows computations on encrypted data without decrypting it.	End-to-end encryption; highly secure.	Extremely computationally intensive; limited applicability to complex models.
Federated Averaging with Weighted Aggregation	Aggregates model updates with weights based on data quantity or quality.	Improves accuracy in non-IID settings; reduces bias in global model.	Requires additional computation for weighting; may not fully address privacy.
Model Compression	Reduces the size of model updates (e.g., quantization, sparsification).	Reduces communication overhead; improves scalability.	May lose some information; requires careful tuning to avoid accuracy loss.
Byzantine-Resilient Aggregation	Filters out malicious or unreliable updates during aggregation.	Improves robustness against adversarial attacks.	May exclude legitimate updates; adds complexity to the aggregation process.
Knowledge Distillation	Shares knowledge (e.g., soft labels) instead of raw data or model parameters.	Reduces communication costs; preserves privacy.	Requires a teacher model; may lose fine-grained details.

TECHNIQUES TO MITIGATE THE IMPACTS OF PRIVACY-PRESERVING TECHNIQUES ON MODEL ACCURACY

As discussed in previously, even though privacy-preserving techniques are crucial for ensuring data security in Federated Learning (FL), they can sometimes hinder model accuracy. However, various

strategies have been developed to minimize these impacts, allowing FL systems to achieve a balance between strong privacy guarantees and high model performance.

Adaptive Noise Addition

Differential privacy is a widely used technique in FL to ensure that individual data points cannot be inferred from the model updates. However, adding noise to the gradients or model parameters can reduce the accuracy of the global model. Adaptive differential privacy addresses this issue by dynamically adjusting the amount of noise added based on the sensitivity of the data and the training progress [13].

Noise Scaling: During the initial stages of training, when the model is far from convergence, higher levels of noise can be added without significantly impacting accuracy. As the model approaches convergence, the noise level can be reduced to fine-tune the parameters.

Layer-Specific Noise: Different layers of a neural network may have varying sensitivities to noise. Adaptive techniques can apply higher noise levels to less sensitive layers (e.g., early layers) and lower noise levels to critical layers (e.g., output layers), preserving overall model accuracy. For instance, the study in [14] introduces a noise-aware algorithm that adaptively adjusts the noise added to client updates based on local data distribution and noise levels. This method seeks to improve overall model accuracy while ensuring strong differential privacy guarantees.

Hybrid Approaches

In Federated Learning (FL), relying solely on one privacy-preserving method can often result in trade-offs that are difficult to balance. For example, while Differential Privacy (DP) can effectively protect sensitive information by adding noise, it may degrade model performance. Similarly, Secure Multi-Party Computation (SMPC) and Homomorphic Encryption (HE) provide robust security but often suffer from high computational overhead and slower training processes. To address these challenges, researchers and practitioners are exploring hybrid approaches that combine multiple techniques, leveraging their complementary strengths while mitigating their individual drawbacks.

Differential Privacy + Secure Multi-Party Computation (SMPC)

This combination enhances both security and efficiency. By adding noise to model updates, DP ensures that individual data points cannot be reverse-engineered from the shared model gradients. However, DP alone can be susceptible to attacks on the noisy updates during transmission. Encrypting these noisy updates with SMPC prevents adversaries from accessing even the perturbed gradients. SMPC ensures that no single party, including the central server, can view the unencrypted data, creating a double layer of protection. This approach makes it nearly impossible for attackers to extract sensitive information, even if they intercept data during transmission. Encrypting only the noisy updates (instead of raw data) reduces the computational load associated with SMPC, making it more scalable for real-world applications *[15]*.

Noise-Optimization Techniques (Enhanced DP + Gradient Clipping)

While Differential Privacy ensures privacy by adding noise to model updates, excessive noise can reduce model accuracy. A hybrid approach combining gradient clipping with optimized noise calibration helps mitigate this effect.

Gradient Clipping restricts extreme gradient values before applying noise. This ensures that the noise added is proportional to the clipped gradients, reducing unnecessary perturbations and dynamically adjusting the level of noise based on the sensitivity of the gradients and the privacy budget can minimize accuracy loss maintaining the utility of model updates by avoiding excessive distortion, thus preserving accuracy while still adhering to privacy guarantees [16].

Partial Homomorphic Encryption (PHE) with Secure Aggregation

Full-scale Homomorphic Encryption (HE) is computationally intensive and can slow down the training process, indirectly impacting accuracy due to delayed convergence. A hybrid approach combining Partial HE with secure aggregation offers a balance as Partial Homomorphic Encryption encrypts only specific sensitive gradients instead of the entire model, significantly reducing computational overhead and Secure Aggregation allows the server to aggregate encrypted updates securely without decrypting them, preserving privacy during the aggregation process. This method improves computational efficiency, ensuring faster convergence and reducing the likelihood of accuracy degradation.

Federated Averaging with Weighted Aggregation

In FL, model updates from different clients are aggregated to improve the global model. However, clients often have non-independent and identically distributed (non-IID) data, which can lead to biased or inaccurate global models. Weighted aggregation techniques address this issue by assigning higher weights to updates from clients with more representative or higher-quality data [17].

Data Quantity-Based Weighting: Clients with larger datasets contribute more to the global model, as their updates are likely to be more reliable.

Data Quality-Based Weighting: Clients with higher data quality (e.g., fewer missing values or noise) are given higher weights during aggregation.

Performance-Based Weighting: Clients whose local models achieve higher accuracy on validation datasets are prioritized in the aggregation process.

Transfer Learning for Non-IID Data

Federated Learning (FL) faces significant challenges when data is non-IID (non-independent and identically distributed). This issue arises because data across participating institutions or devices often has unique characteristics influenced by local demographics, devices, or environments. In healthcare, for instance, patient data varies significantly across hospitals due to differences in population health, medical equipment, and diagnostic procedures. Such heterogeneity can slow convergence, degrade global model performance, and lead to biased or suboptimal outcomes. To address these challenges, transfer learning techniques have proven invaluable.

Pre-Training on Public Datasets

One effective strategy is leveraging public datasets to pre-train the global model before federated training begins. This process enables the model to learn general features that are transferable to specific tasks, even when local data distributions differ significantly. A global model is first trained on publicly available, large-scale datasets (e.g., ImageNet for imaging tasks, MIMIC-III for electronic health records). These datasets, while not perfectly aligned with specific local datasets, provide the model with foundational knowledge, such as recognizing anatomical structures or understanding common clinical patterns.
Once pre-trained, the model is distributed to local nodes, where it is fine-tuned on institution-specific data during federated training [18].

Knowledge Distillation

Another promising approach is knowledge distillation, where local models share distilled knowledge with the global model to enhance generalization across non-IID datasets. Instead of directly sharing model parameters or gradients, local models transmit distilled information, such as soft labels (probabilistic outputs) or feature representations derived from their training data. The global model aggregates this distilled knowledge to improve its performance while maintaining privacy and reducing communication overhead [19].
The global model, enhanced by the distilled knowledge, updates its structure to better generalize across diverse local datasets. This process avoids overfitting to dominant datasets and ensures equitable

performance across institutions. This enables cross-domain learning by allowing the global model to incorporate patterns from diverse datasets, such as diagnostic codes from urban and rural hospitals and reduces bias in the global model by emphasizing shared knowledge rather than raw gradients, which may overfit to local peculiarities while enhancing data privacy, as sensitive patient data remains local, and only abstracted knowledge is shared.

CASE STUDIES: FEDERATED LEARNING IN HEALTHCARE

Federated Learning (FL) has been successfully applied in various healthcare domains, demonstrating its potential to enable collaborative research while preserving data privacy. This section explores three prominent case studies that highlight the practical applications of FL in medical imaging, electronic health records (EHR), and genomics. Each case study provides valuable insights into how FL addresses the challenges of data privacy and model accuracy in real-world healthcare scenarios.

A groundbreaking study by Sheller et al. [20] demonstrated the use of FL for diagnostics. The study involved collaboration among 10 healthcare organizations, each contributing locally stored MRI scans of brain tumor patients. The FL model achieved comparable accuracy to a centrally trained model, with a Dice similarity coefficient (DSC) of 0.85, indicating high segmentation precision.

Data privacy was maintained throughout the process, as only model updates (gradients) were shared with the central server, and no raw MRI scans were exchanged. This study demonstrated that FL could enable large-scale collaboration in medical imaging research, overcoming data silos and privacy concerns. Table II summarizes several case studies of FL in healthcare.

COMPARATIVE ANALYSIS: MACHINE LEARNING VS. FEDERATED LEARNING IN HEALTHCARE

The adoption of machine learning (ML) and federated learning (FL) in healthcare has revolutionized the way data is utilized for diagnostics, treatment, and research. However, the two approaches differ significantly in their handling of accuracy and privacy. This section provides a comparative analysis of ML and FL in healthcare, focusing on their trade-offs between accuracy and privacy, supported by genuine resources and studies.

Traditional ML relies on centralized data processing, where all data is pooled into a single location for model training. Centralized ML models often achieve high accuracy because they have access to the entire dataset, enabling the model to learn from a comprehensive and diverse set of examples but requires the sharing of raw patient data, which poses significant privacy risks. Data breaches in healthcare can lead to the exposure of sensitive information, such as medical histories and genetic data for example, the 2015 Anthem data breach exposed the personal information of nearly 78.8 million individuals, highlighting the vulnerabilities of centralized data storage.

TABLE 2: CASE STUDIES OF FEDERATED LEARNING USED IN HEALTHCARE

Case Study	Application	FL Methodology	Results	Impact
Multi-Center Critical Care Research [21]	Patient survival prediction in ICUs	Federated Averaging	Achieved comparable performance to centralized models while maintaining strict data privacy.	Enabled multi-institutional collaboration without sharing sensitive patient data, proving FL's viability for critical care applications.
Collaborative Federated Learning for COVID-19 Diagnosis at the Edge[22]	COVID-19 diagnosis using multi-modal data	Edge-based Federated Learning	Improved diagnostic accuracy by integrating diverse data sources.	Demonstrated FL's ability to enhance access to diagnostic tools in remote and resource-constrained healthcare settings while preserving privacy.
Decentralized FL for Cancer Classification [23]	Cancer classification using medical images	Federated Averaging (FedAvg) and Federated Proximal (FedProx) algorithms	Both algorithms effectively handled non-IID data distributions, achieving high classification accuracy across cervical, lung, and colon cancer datasets	Demonstrated the potential of decentralized FL to maintain data privacy while achieving accurate cancer classification across multiple institutions.

Overall, centralized ML is often infeasible in healthcare due to strict privacy regulations like HIPAA and GDPR, which restrict the sharing of patient data across institutions and data silos created by these regulations limit the ability to train robust models, particularly in rare diseases or specialized treatments where data is scarce.

However, as we have discussed earlier FL addresses the privacy limitations of centralized ML by enabling collaborative model training without sharing raw data. Instead, only model updates (e.g., gradients) are shared, preserving data privacy.

FL provides strong privacy guarantees by keeping data localized and using techniques like differential privacy, secure multi-party computation (SMPC), and homomorphic encryption. For instance, Kaissis et al. (2020) used SMPC in FL for genomic research, ensuring that no single institution could access raw genetic data while achieving 92% classification accuracy [24] and FL enables cross-institutional collaboration without violating privacy regulations, making it feasible to train models on diverse datasets.

CONCLUSION

Federated Learning represents a transformative approach to collaborative machine learning in healthcare, offering a way to leverage distributed data while preserving privacy. However, the privacy-accuracy trade-off remains a significant challenge that requires careful consideration. By exploring recent advancements, case studies, and future directions, this article highlights the potential of FL to revolutionize healthcare while addressing the critical issue of data privacy. As the field continues to evolve, it is essential for researchers and practitioners to work together to overcome the technical, regulatory, and ethical challenges associated with FL in healthcare.

ACKNOWLEDGMENT

REFERENCES

Anand, A., 2023. GDPR and Healthcare: Balancing Data Privacy and Access to Medical Information. NUJS J. Regul. Stud., 8, p.27.

Bhosale, K.S., Nenova, M. and Iliev, G., 2021, September. A study of cyber attacks: In the healthcare sector. In 2021 Sixth Junior Conference on Lighting (Lighting) (pp. 1-6). IEEE.

Kalapaaking, A.P., Stephanie, V., Khalil, I., Atiquzzaman, M., Yi, X. and Almashor, M., 2022. Smpc-based federated learning for 6g-enabled internet of medical things. IEEE Network, 36(4), pp.182-189.

Khalid, N., Qayyum, A., Bilal, M., Al-Fuqaha, A. and Qadir, J., 2023. Privacy-preserving artificial intelligence in healthcare: Techniques and applications. Computers in Biology and Medicine, 158, p.106848.

Li, Tian, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. "Federated learning: Challenges, methods, and future directions." IEEE signal processing magazine 37, no. 3 (2020): 50-60.

Liu, H., Li, C., Liu, B., Wang, P., Ge, S. and Wang, W., 2021, December. Differentially private learning with grouped gradient clipping. In Proceedings of the 3rd ACM International Conference on Multimedia in Asia (pp. 1-7).

Mbonihankuye, S., Nkunzimana, A. and Ndagijimana, A. 2019 'Healthcare Data Security Technology: HIPAA compliance', Wireless Communications and Mobile Computing, 2019, pp. 1–7. doi:10.1155/2019/1927495.

Mehrjou, A., Soleymani, A., Buchholz, A., Hetzel, J., Schwab, P. and Bauer, S., 2022. Federated learning in multi-center critical care research: A systematic case study using the eicu database. arXiv preprint arXiv:2204.09328.

Qayyum, A., Ahmad, K., Ahsan, M.A., Al-Fuqaha, A. and Qadir, J., 2022. Collaborative federated learning for healthcare: Multi-modal covid-19 diagnosis at the edge. IEEE Open Journal of the Computer Society, 3, pp.172-184.

Reuters. (2015). Anthem Hacking Points to Security Challenges of Healthcare Data. Retrieved from https://www.reuters.com/article/us-anthem-cybersecurity

Seh AH, Zarour M, Alenezi M, Sarkar AK, Agrawal A, Kumar R, Khan RA. Healthcare Data Breaches: Insights and Implications. Healthcare (Basel). 2020 May 13;8(2):133. doi: 10.3390/healthcare8020133..

Sheller, M.J., Edwards, B., Reina, G.A., Martin, J., Pati, S., Kotrotsou, A., Milchenko, M., Xu, W., Marcus, D., Colen, R.R. and Bakas, S., 2020. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. Scientific reports, 10(1), p.12598.

Staynings, Richard. "Cybersecurity." In Digital Health Entrepreneurship, pp. 131-155. Cham: Springer International Publishing, 2023.

Subramanian, M., Rajasekar, V., VE, S., Shanmugavadivel, K. and Nandhini, P.S., 2022. Effectiveness of decentralized federated learning algorithms in healthcare: a case study on cancer classification. Electronics, 11(24), p.4117.

Sun, T., Li, D. and Wang, B., 2022. Decentralized federated averaging. IEEE Transactions on Pattern Analysis and Machine Intelligence, 45(4), pp.4289-4301.

Tan, Y., Long, G., Ma, J., Liu, L., Zhou, T. and Jiang, J., 2022. Federated learning from pre-trained models: A contrastive learning approach. Advances in neural information processing systems, 35, pp.19332-19344.

Truex, S., Baracaldo, N., Anwar, A., Steinke, T., Ludwig, H., Zhang, R. and Zhou, Y., 2019, November. A hybrid approach to privacy-preserving federated learning. In Proceedings of the 12th ACM workshop on artificial intelligence and security (pp. 1-11).

Wei, Kang, et al. "Federated learning with differential privacy: Algorithms and performance analysis." IEEE transactions on information forensics and security 15 (2020): 3454-3469.

Wen, J., Zhang, Z., Lan, Y., Cui, Z., Cai, J., & Zhang, W. (2023). A survey on federated learning: challenges and applications. International Journal of Machine Learning and Cybernetics, 14(2), 513-535.

Wu, C., Wu, F., Lyu, L., Huang, Y. and Xie, X., 2022. Communication-efficient federated learning via knowledge distillation. Nature communications, 13(1), p.2032.

Yang, X., Huang, W., & Ye, M. (2023). Dynamic personalized federated learning with adaptive differential privacy. Advances in Neural Information Processing Systems, 36, 72181-72192.

Yang, Z., Zhou, M., Yu, H., Sinnott, R.O. and Liu, H., 2022. Efficient and secure federated learning with verifiable weighted average aggregation. IEEE Transactions on Network Science and Engineering, 10(1), pp.205-222.

Zhang, Jinghui, Dingyang Lv, Qiangsheng Dai, Fa Xin, and Fang Dong. "Noise-aware local model training mechanism for federated learning." ACM Transactions on Intelligent Systems and Technology 14, no. 4 (2023): 1-22.

Zhu, Hangyu, Jinjin Xu, Shiqing Liu, and Yaochu Jin. "Federated learning on non-IID data: A survey." Neurocomputing 465 (2021): 371-390.