National Centre of Excellence
CYBERSECURITY TECHNOLOGY AND ENTREPRENEURSHIP

इलेक्ट्रॉनिकी एवं सूचना प्रौद्योगिकी मंत्रालय
MINISTRY OF ELECTRONICS AND INFORMATION TECHNOLOGY
सत्यमेव जयते

DSCI
PROMOTING DATA PROTECTION
A nasscom Initiative

Exploratory Note on

FEDERATED LEARNING IN HEALTHCARE SECURITY DOMAIN

Objective

A high-level summary of the security issues, legislative and regulatory changes, and industry best practices related to federated learning in healthcare applications is given in this note. This document identifies important factors to take into account for any policy interventions in this area, even though it does not offer specific solutions to the challenges of securing federated learning systems. Future studies and technological developments aiming at enhancing the security and privacy of federated learning in healthcare may be guided by the ideas discussed in the last section.

Contributors

Sonam Lata, Assistant Professor
Department of Computer Science and Engineering,
IILM University Gurugram, Haryana India

Shabana Mehfuz, Professor
Department of Electrical Engineering
Jamia Millia Islamia

Table of CONTENTS

I Context 4

II Introduction 5

III Federated Learning Technique 8

IV Latest Development Applications 15

V Challenges of Federated Learning 19

VI Possible Future Research Directions 23

VII Conclusion 25

1

Context

A renewed interest in improving accessibility of medical data has been spurred by recent developments in deep learning for healthcare and computer-aided laboratory services. Due to data-driven medical research and the growing digitization of medical records, data privacy has emerged as a significant issue in the healthcare industry. Sensitive patient data must be protected from breaches and unauthorized access since they may result in serious ethical and legal ramifications.

Because deep learning models are specifically made to embrace a data-driven approach, they have demonstrated exceptional performance in this area. Exposure to larger datasets is beneficial for these models as it allows them to continuously improve their performance.

However, worries about data ownership, privacy, and legal constraints have surfaced as healthcare organizations work to compile clinical records onto central servers in order to build strong deep learning models. It is a difficult balancing act to protect private medical information while utilizing the collective expertise of several healthcare facilities. Using privacy-preserving techniques, which enable the use of data from multiple centres without compromising security, is one promising way to address these concerns. Large machine learning models that have been trained across several data centres can now be deployed thanks to a technique called federated learning (FL), which eliminates the need to share sensitive data. The three most innovative and in-demand technologies in the field of intelligent healthcare are federated learning (FL), artificial intelligence (AI), and explainable artificial intelligence (XAI). The way the healthcare system has historically operated has involved centralized agents exchanging unprocessed data. Thus, there are still significant flaws and difficulties with this system. When AI is integrated, the system would consist of several agent collaborators that can effectively communicate with their intended host. Once more, FL is an intriguing feature that operates decentralized and keeps the model-based communication going in the chosen system without sending the raw data. The integration of FL, AI, and XAI methodologies has an opportunity to reduce different constraints and difficulties encountered within the healthcare system

1

Introduction

Practically speaking, data is usually distributed rather than centralized and needs to be collected beforehand from a variety of sources. Moreover, the information that an organization acquires may have an odd distribution that makes it difficult to develop systematic models. For example, hospitals across different regions or countries appear to have different patient profiles and pathologies; therefore, they should work together to design and develop ML applications that are equally beneficial to all of their patients.

Overview of Federated Learning

A machine learning (ML) framework called Federated Learning (FL) allows clients located in different parts of the world to work together to train a common model without directly sharing local data. Class imbalances and non-IID (non-independent and identically distributed) data distributions among clients are among the issues this framework addresses. FL enhances training of models while preserving data security and privacy in distributed environments through the regular distribution of updated models and the use of techniques such as data harmonization and adaptive self-distillation.

FL is gaining traction in the healthcare industry because of its use in data privacy, which is crucial when managing sensitive patient medical data, including reports, pictures, electronic health records (EHRs), and so forth. Together, FL algorithms and healthcare organizations can train ML models while preserving the security and localization of raw data, protecting patient privacy. This method has been used in many different areas of healthcare, especially in the diagnosis of illnesses. The accuracy of diagnosing diseases like cancer, diabetic retinopathy, and other ailments has increased thanks to models trained on a variety of multi-institutional datasets.

Figure 1 illustrates the FL framework used to train various TL-based models, such as VGG16, ResNet50, and ResNet101, across two separate client servers.

The trained weights that each client exchanges are used as model parameters by the global model. Furthermore, FL has made it possible for EHR data to be shared securely, promoting the creation of patient outcome prediction models and individualized treatment plans while upholding privacy standards. FL facilitates

inter-institutional data collaboration in drug discovery, which enhances the identification of putative therapeutic compounds. Additionally, FL has made it easier to analyse data from wearable devices, which helps with the early diagnosis and tracking of chronic illnesses like diabetes and heart disease. FL promotes data collaboration while safeguarding private medical information by placing a strong emphasis on security and privacy.

Importance of Data privacy in healthcare

The amount of digital healthcare data has increased dramatically in recent years. Studies using substantial volumes of data gathered from diverse sources are frequently conducted in clinical research. Medical data is accessible to the pharmaceutical industry, individuals, insurance companies, and health institutions. Moreover, every institution might be associated with a distinct group of stakeholders. These data are sensitive and are not accessible. In order to preserve patient privacy, it is morally and legally necessary to assemble and transmit these datasets. New laws that regulate data sharing while protecting user security and privacy have been passed by the majority of healthcare facilities, national laws and regulatory bodies such as the Health Insurance Portability and Accountability Act (HIPAA) and the General Data Protection Regulation (GDPR). Furthermore, a key component of patient rights is access to information and control over the transmission, storage, and use of medical data.

Due to the sensitive and dispersed nature of EHRs (Electronic Health Records) in real-world settings, an efficient method for learning from data stored in hospitals and other health- related institutions while protecting data privacy is required. This encourages us to investigate federated learning's potential and benefits for the healthcare industry.

Key Security Challenges Addressed by FL in Healthcare

The following noteworthy events over the years have had a big impact on how FL is used in healthcare:

- **Privacy-Preserving Data Sharing:** Without disclosing raw patient data, FL allows healthcare organizations to work together to develop strong ML models. This is especially crucial for sensitive medical data, like EHRs, and picture data. FL preserves patient privacy by utilizing a larger, more varied dataset to enhance model accuracy while keeping data localized and training models across several institutions.

- **Enhancing Disease Diagnosis and Detection:** Federated models trained on diverse datasets from multiple healthcare institutions can detect diseases like cancer, diabetic retinopathy, and COVID-19 more accurately. They benefit from the shared expertise and data diversity across institutions. This collaborative approach leads to the early detection and better diagnosis of diseases, improving patient outcomes by providing clinicians with more reliable diagnostic tools.
- **Enhancing Predictive Models for Patient Outcomes:** FL encourages the creation of predictive models that examine patient medical records and treatment schedules in order to predict outcomes and customize care. These predictive models can be used by clinics and hospitals to identify high-risk patients, make proactive adjustments to treatment plans, and lower the rate of readmissions.
- **Drug Development and Discovery:** Without disclosing confidential information, FL allows research organizations and pharmaceutical companies to work together to analyze data about the safety and efficacy of drugs. This methodology expedites the process of identifying putative therapeutic compounds and facilitates a more thorough evaluation of drug effects among heterogeneous patient populations.
- **Wearable Device Data Analysis:** FL makes it possible to analyze continuous data streams decentralized, which aids physicians in tracking and forecasting trends in chronic illness. Better patient management results from greatly increased early detection and monitoring of diseases like diabetes and heart disease.

Federated Learning Technique

Horizontal Federated Learning

Horizontal Federated Learning (HFL) represents a distributed ML paradigm aimed at enhancing model performance by harnessing data from decentralized devices or nodes. This method, which is frequently used in situations involving mobile or Internet of Things (IoT) devices, allows for collaborative learning while maintaining data security and privacy. Horizontal federation divides data horizontally across several nodes so that each node can process its share of the data locally. This decentralized method addresses serious privacy and data security concerns by lowering the risk of data breaches and guaranteeing that sensitive data stays on users' devices as shown in figure 1.

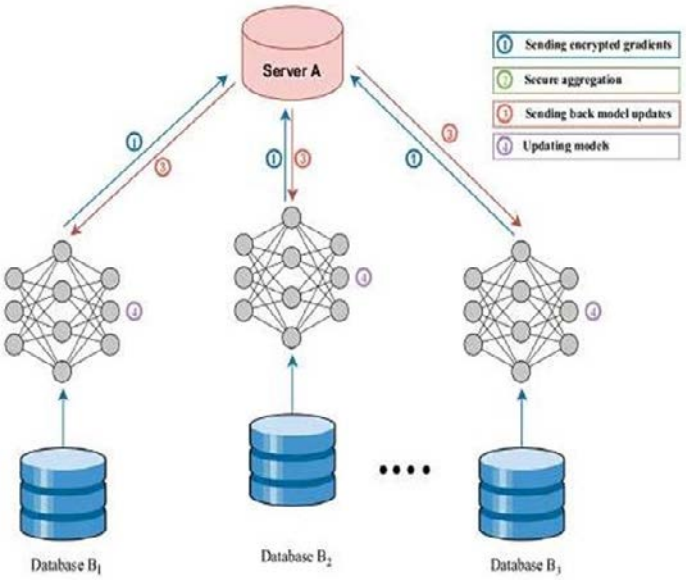

Figure 1: Horizontal FL in the healthcare sector [1]

One significant benefit of HFL over centralized ML techniques is its capacity to utilize a bigger and more varied dataset. HFL makes it possible for models to be trained on a wide variety of data sources, improving accuracy and robustness, by aggregating data from several nodes. Additionally, there may be advantages to this distributed learning strategy in the healthcare industry, as it

permits cooperative training on electronic health records (EHR) from several institutions without jeopardizing patient privacy as shown in Table 1.

Table 1: Application and related research work

Application	Related Research
Heart Disease Prediction	Hybrid Classifier-Based FL in Health Service Providers for Cardiovascular Disease Prediction.[2]
Mortality Prediction	FL of electronic health records to improve mortality prediction in hospitalized patients with COVID-19: machine learning approach.[3]
MRI Image Analysis	FL of generative image priors for MRI reconstruction[4]
CT Image Analysis	Blockchain-federated learning and deep learning models for covid-19 detection using ct imaging.[5]
fMRI Image Analysis	Multi-site fMRI analysis using privacy-preserving FL and domain adaptation: ABIDE results.[6]
COVID-19 Detection	Dynamic-fusion-based federated learning for COVID-19 detection.[7]
Skin Cancer Detection	Multimodal melanoma detection with FL.[8]
Diabetic Retinopathy Decision Support	Making FL robust to adversarial attacks by learning data and model association.[9]
Heart Failure Decision Support	FedSDM: FL based smart decision making module for ECG data in IoT integrated Edge-Fog-Cloud computing environments.[10]
Electronic Health Record Text Analysis	Protecting personal healthcare record using blockchain & FL technologies.[11]

However, despite its advantages, HFL presents several challenges that need to be addressed. Ensuring data consistency and quality across distributed nodes is a critical concern, as biases and errors in individual node data can significantly impact model performance. Strategies such as data pre-processing and quality checks are employed to mitigate these challenges and ensure the reliability of the training process.

Another major challenge in HFL is continuing to implement the training process in an efficient and scalable manner. In order to handle the resource-intensive nature of the task, training on combined data from multiple nodes requires optimization and scaling strategies, such as distributed training and parallel processing. HFL presents a viable way to increase model accuracy in distributed environments while maintaining privacy in spite of these difficulties. Furthermore, Table 1 shows the various uses of HFL in the medical field, especially in predictive modelling. The table presents key findings from a range of studies, offering a thorough summary of the viewpoints and approaches taken by authors in integrating HFL into healthcare systems.

Vertical Federated Learning

Vertical Federated Learning (VFL) aims to enhance the efficacy of ML models by leveraging data from diverse distributed nodes or devices. This method involves vertically partitioning the data of each node, ensuring that each node possesses a distinct set of attributes or characteristics, and subsequently training the model using collective data from all nodes. The proposed approach aligns with findings. Applications of VFL are widely used in industries like finance and healthcare, where heterogeneous data is frequently collected from various devices and nodes. VFL makes it possible to aggregate patient medical records, providing distinct data for model training in the healthcare industry. In contrast to centralized ML techniques, this study suggests a hybrid ML approach for medical diagnosis that expands the dataset substantially by training the model on data from all patients. In Figure 2, the use of VFL in the healthcare system is shown graphically.

Vertical data partitioning, which separates data from each node into discrete subsets corresponding to particular data categories, is a step in the pre-processing stage of VFL. The data is then added to the model during training, which improves overall performance by allowing the model to learn from a bigger and more diverse dataset.

One major challenge in VFL is ensuring high-quality performance and consistency across all nodes. VFL usually uses methods for data refinement and quality control to overcome this obstacle. One acknowledged difficulty with VFL training is its scalability and practicality. To

address these issues, optimization and scaling techniques like distributed training and parallel processing are frequently used. Equations for VFL that are derived from the optimization problem of minimizing the loss function while maintaining data privacy are presented in the study.

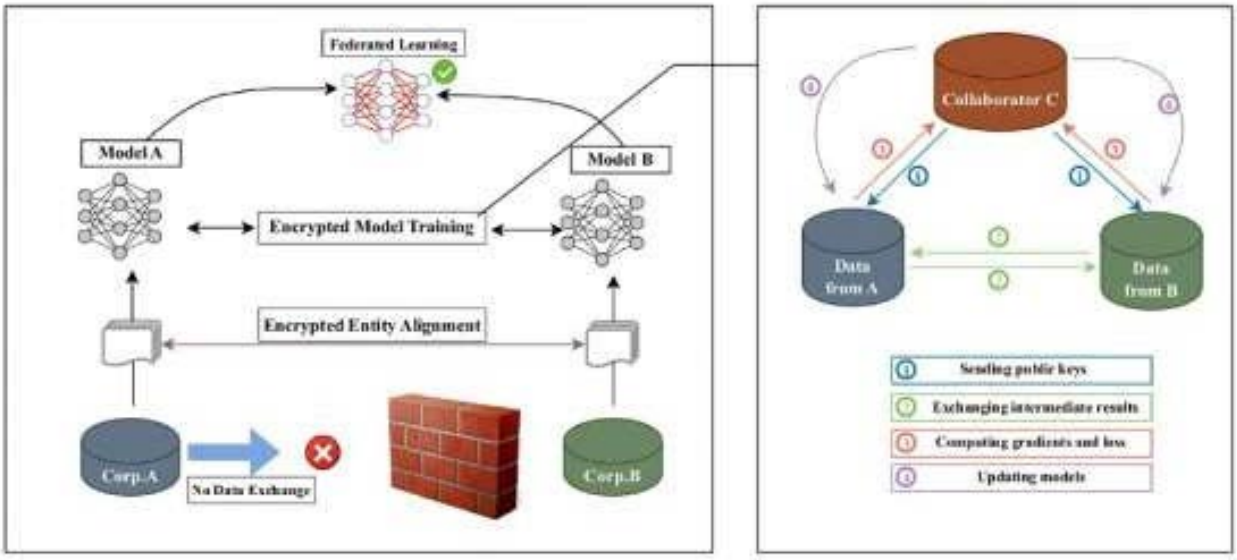

Figure 2: Vertical FL in the healthcare sector [2]

Table 2: Vertical FL and its uses in the Healthcare Sector

Disease Diagnosis	Heart Disease Diagnosis
	Cancer Diagnosis
	COVID-19 Diagnosis
Clinical Decision Support	Predicting Adverse Drug Reactions
	Diagnosis Support for Rare Diseases
Electronic Heath Record Analysis	Clinical Outcome Predictions
Imaging Analysis	MRI Image Segmentation
	PET Image Analysis

Using collective data from decentralized devices or nodes, Vertical FL is a distributed machine learning technique that protects user privacy and allows for the creation of more accurate models. Compared to traditional centralized machine learning techniques, this methodology enables training on a much larger and more diverse dataset, with potentially exciting implications across a range of industries, including finance and healthcare. Table 2 provides insights from several studies on the subject in addition to demonstrating the many uses of VFL in the field of healthcare through its predictive modelling capabilities.

Federated Transfer Learning

An ML technique called Federated Transfer Learning (FTL) aims to protect data privacy as it moves from a centralized to a decentralized model. Fan et al.'s work on IoT Defender

proposes a method wherein a centralized model is trained initially, and then trained information is transferred to decentralized models that are trained on smaller datasets. Through the assimilation of information from the centralized model, that approach seeks to improve the accuracy and performance of distributed models.

FTL preserves the privacy of individual data while enhancing model performance in new domains by utilizing knowledge gained from one to another. In situations where sensitive or dispersed data is involved, this strategy is especially pertinent as it reduces the possibility of data breaches and unwanted access. In IoT and mobile device environments, FTL application has become commonplace, addressing the constraints posed by dispersed nodes or devices with constrained resources.

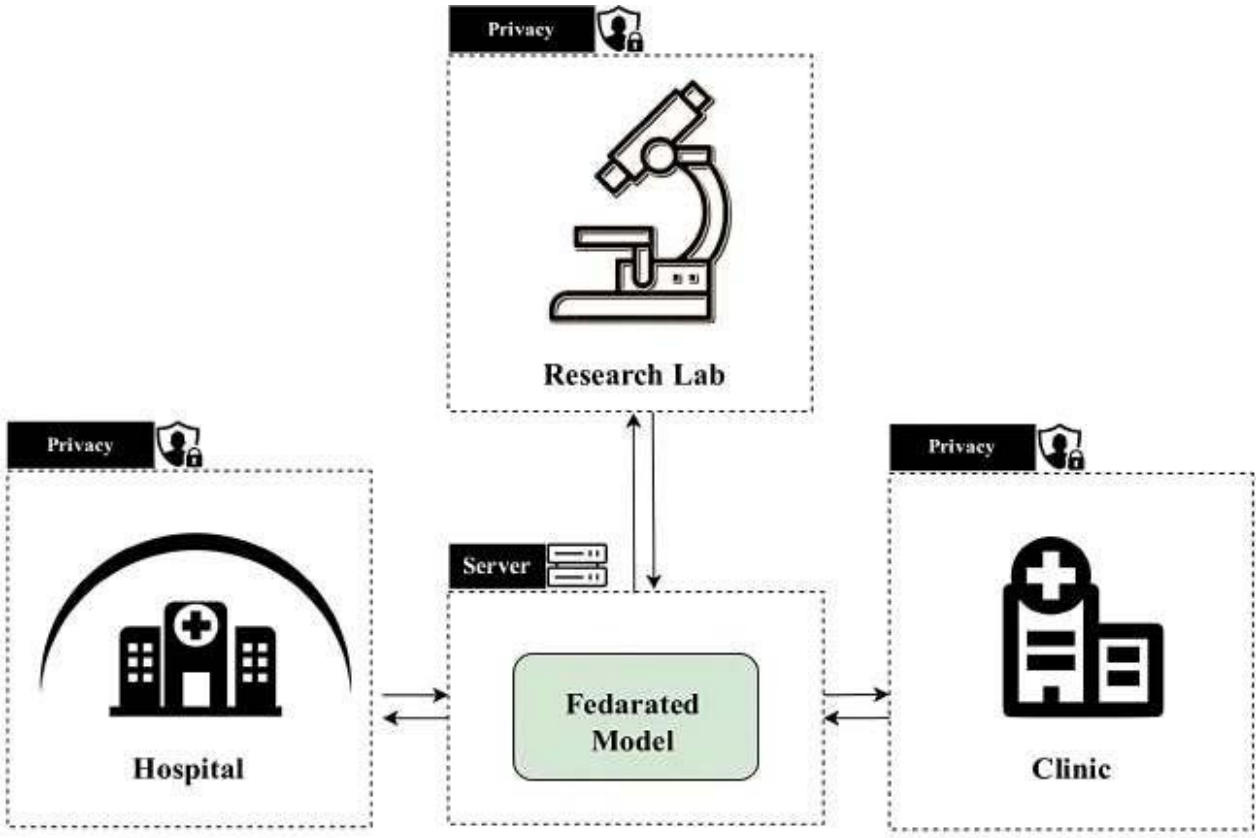

Figure 3: Federated Transfer Learning in the healthcare sector.[13]

This method does away with the requirement for centralized data storage, enabling safe and private learning processes. Zheng, Xiao, et al. used techniques like fine-tuning, knowledge distillation, and model compression to investigate the effectiveness of FTL in knowledge transfer across domains.

In additional research, the authors investigate how knowledge distillation and model compression both of which use the central model—relate to one another in order to support local training as well as distributed model training. Furthermore, Zhang, Wei, et al. explore the mathematical formulations governing knowledge transfer in federated learning with a thorough analysis of equations for FTL.

Their analysis underscores the importance of problem-specific settings and the nature of data in FTL equations. The use of FTL in the healthcare industry is graphically depicted in Figure 3, which also offers insights into the functionality and operation of the technology.

Federated Domain Adaption

One well-known machine learning technique that is frequently used in a variety of domains is federated domain adaptation, or FDA. FDA addresses the issue of domain shift and protects user privacy while promoting knowledge transfer across domains. The approach has been well-documented in earlier studies, showing encouraging outcomes, especially in smart systems. FDA aligns distributions between a target domain, which is frequently characterized by limited data and annotations, and a source domain, which is usually data-rich, by utilizing domain adaptation algorithms. The FDA improves model generalizability by lowering the divergence between domains, particularly in situations where target domain data is sparse. The FDA's

use in the healthcare industry is graphically depicted in Figure 4, which also offers insights into the organization's operations in this field. FDA facilitates the development of accurate models while maintaining data confidentiality and is used in situations where data centralization is either impractical or undesirable. In order to train models appropriate for target domains with sparse data and different distributions from source domains, this method carefully makes use of the data that already exists. The adaptation process involves adjusting a pre-trained model from the source domain to fit the target domain's data distribution, ensuring sensitive data remains on users' devices. After that, a global model customized for the target domain is produced by combining the updated model parameters. The FDA benefits scenarios with limited target domain data and distinct distributions by facilitating knowledge transfer across domains while maintaining privacy.

Multitask Federated Learning

With the use of distributed data, Multitask Federated Learning (MTFL) allows the training of multiple models in parallel, each of which is devoted to a distinct task. These models work together to share knowledge and enhance overall performance as a group. When data is scattered across multiple devices or nodes and there are interrelated activities involved, MTFL has become more popular. Through the use of contextual knowledge from various tasks, MTFL improves system performance and model training accuracy. This method is especially helpful in situations where decentralization of data is required or desired. Figure 4 and Figure 5 provides an example of how MTFL is used in the healthcare industry and highlights some of its features.

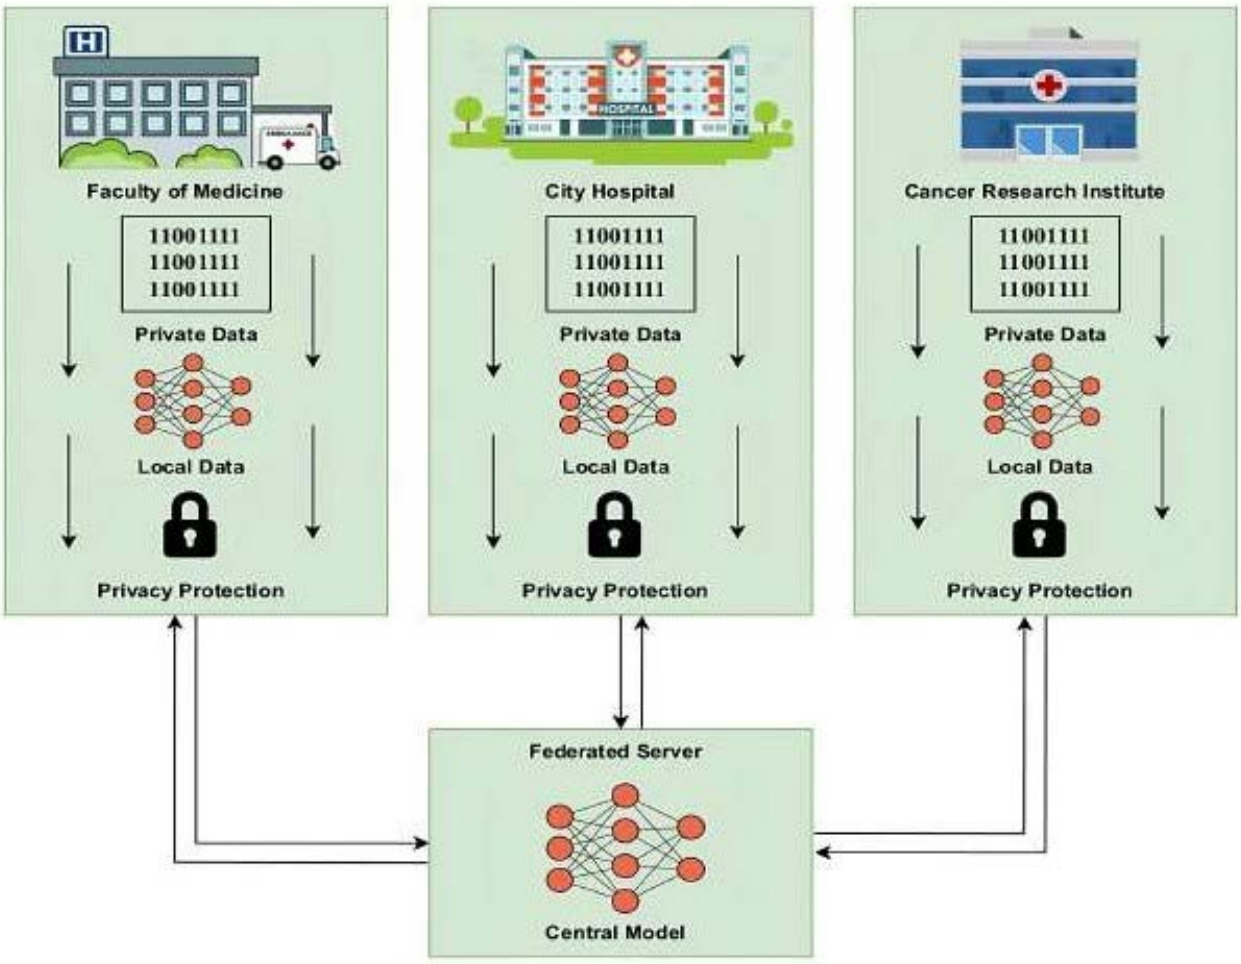

Figure 4: Federated domain adaption in the healthcare sector.[14]

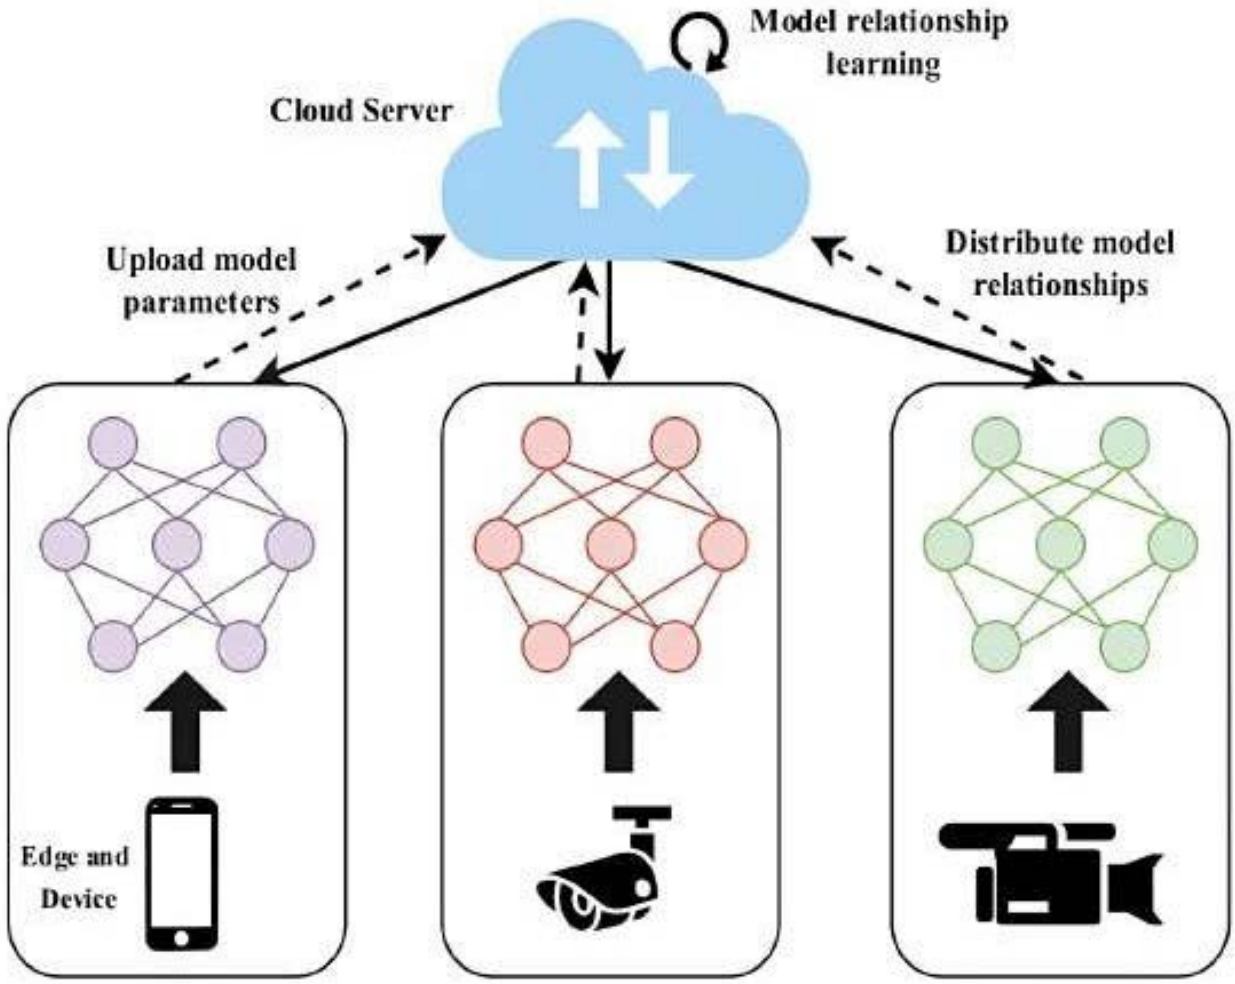

Figure 5: Multitask FL in the healthcare sector.[15]

Iterative model training is done in MTFL, where local parameter updates are aggregated at a central node. This allows each device's model to benefit from aggregated parameters, improving training efficacy. Achieving high accuracy across tasks, MTFL maintains privacy and security while optimizing joint learning across multiple tasks in a federated setting. With global parameters updated iteratively, the objective function minimizes the total loss of functions across tasks and devices.

Federated Meta Learning

Federated Meta-Learning (FML) leverages prior knowledge from a variety of tasks and decentralized data sources to build models that can quickly adapt to new challenges. In FML, models are trained iteratively on several tasks, revising their parameters in light of new information, and then applying the modified models to new tasks. This method saves time and resources by allowing for a quicker adaptation to new tasks without requiring complete retraining. Furthermore, FML solves distributed data issues by allowing model training on several devices without the need for centralized data repositories. FML simultaneously optimizes meta-models on different devices, using shared meta-models to learn about tasks and fine-tuning on local tasks. This optimization ensures data confidentiality, enables quick adaptation to new tasks with little data, and reduces meta-model loss across validation tasks.

Furthermore, FML effectively aggregates decentralized data to train task-agnostic models that are versatile and adaptable, enabling them to function well in a variety of tasks or domains.

IV

Latest Development Applications

COVID-19 Detection

FL has been effectively employed in the context of COVID-19 to aid in virus detection using wearable technology and various medical tools. To help with early virus detection, smart watches and activity trackers, for example, can monitor COVID-19 symptoms by collecting data on temperature, heart rate, and other vital indicators. But gathering this much data from millions of devices around the world presents difficulties and concerns about centralized databases, especially with regard to security and privacy. FL allows for the local storage of data on devices and the decentralized training of models in order to address these issues. This method shows promise in protecting patient privacy while using COVID-19 EHR data to create precise predictive models.

An additional use is the identification of COVID-19 via X-ray imaging, a diagnostic technique frequently employed for respiratory illnesses. Salam et al. demonstrated the potential of FL in this domain by creating ML models, including a FL model, using chest X- ray (CXR) images from COVID-19 patients. FL allows for model training on a large dataset while maintaining privacy and security by storing the images locally at hospitals and clinics.

Additionally, FL can be used to detect COVID-19 through CT scans. To do this, models can be trained on a sizable dataset of CT images to recognize unique patterns linked to the virus. By combining blockchain technology and FL for joint model training, Rajesh et al.'s approach ensures anonymity while improving the accuracy of COVID-19 detection. All things considered, FL provides a useful way to improve COVID-19 detection speed and accuracy while maintaining data security and confidentiality in the face of growing wearable and medical equipment use.

Breast Density Classification

In order to determine a patient's risk of developing breast cancer, breast density must be classified. FL, a cutting-edge technology, provides a viable way to improve breast density classification. With FL, several parties can collaborate decentralized while maintaining the confidentiality and privacy of their data. In order to train a machine learning model for breast density classification, a number of hospitals or imaging centres may combine their data while maintaining the anonymity of the individual datasets. One of the primary challenges to breast density classification is the absence of large, high-quality datasets that sufficiently represent the diverse community of women who use mammography. The open-source program MammoDL analyses the density and complexity of breast tissue from mammograms quantitatively using a FL method and the U-Net DL architecture. FL can help to overcome this challenge by merging smaller datasets from multiple institutions into a larger, more diverse dataset. The performance of the machine learning model can be improved by this additional diversity, which can also reduce the risk of over fitting, which occurs when a model is trained too closely on a specific dataset and finds it difficult to generalize to new data. FL might also be helpful with privacy-related concerns.

In a typical machine learning scenario, all data must be centralized, which leaves it vulnerable to hackers and data breaches. In Florida, raw data is only used to update models; it is still kept on the servers of the individual universities. There is less chance of data leakage and patient privacy is protected.

Healthcare Monitoring

By facilitating collaboration between academics and healthcare professionals, FL promises to have a revolutionary effect on healthcare monitoring. Through the consolidation of data from various sources, such as wearable technology, medical devices, and patient records, FL makes it easier to create precise machine learning models that predict patient outcomes. Using information from medical devices, electronic health records, and patient-generated data, these models forecast medication responses, readmission risks, and disease susceptibility. Fitness trackers and other wearable have the ability to remotely monitor vital signs, activity levels, and sleep patterns, notifying healthcare professionals of any abnormalities. Suggested frameworks, such as edge-assisted data analytics, use FL to protect privacy while locally retraining ML models on user-generated data. By keeping data locally and encrypting it before sending it to a central server, FL protects patient data privacy by lowering the possibility of data breaches and unauthorized access. All things considered, FL

has the ability to completely transform healthcare monitoring by enhancing forecast accuracy, offering individualized care, and protecting patient privacy.

Medical Imaging

In FL, medical images are processed decentralized using machine learning techniques, specifically in the field of medical imaging. Medical imagery is essential for patient diagnosis and treatment, which makes this approach critical to the healthcare industry.

Typically, different sources with distinct datasets—such as imaging centers, hospitals, and clinics—are used to gather medical imaging data. As more data becomes available, FL enables these entities to work together to train global models, improving disease identification and treatment efficacy over time. The privacy issues associated with centralized data systems are mitigated by FL, which maintains data security and privacy while utilizing the combined expertise of various organizations to improve model development. FL algorithms, like Federated Averaging (FedAvg), facilitate instantaneous cooperation between entities, expediting the creation of models and the diagnosis of

illnesses. By combining data from various sources, FL also solves data imbalance problems that are frequently encountered in medical imaging, leading to more reliable models and precise diagnoses. In general, FL promises to lower healthcare costs and enhance patient outcomes in applications involving medical imaging. As indicated in Table 3, a number of publications have examined the use of FL in medical contexts and summarized their conclusions, methods, and results.

For example, Joynab et al.[16] use CNN-based FL architectures to detect cervical cancer while balancing data privacy with precise image classification in three different experimental settings. By combining updates from locally trained models into a global model, the suggested system achieves test accuracies of 78.4% in a non-IID setting and 94.36% in an IID setting. The model's performance does, however, differ substantially in IID and non-IID settings, indicating possible difficulties with heterogeneous data. The reliance on FL also requires robust infrastructure and cooperation among institutions, which may not always be feasible.

Privacy concerns in AI training for cancer image analysis are addressed by Truhn et al. [17] using Somewhat-Homomorphically-Encrypted Federated Learning (SHEFL). This technique prevents data breaches during model updates by transferring only encrypted weights. Using multicentre datasets, SHEFL was successfully applied to a variety of cancer image analysis tasks.

Table 3: Summary of FL publications applied in medical applications.

ML Task	Clinical Tasks	Medical Input Data	Model Architecture
Classification	COVID-19 Diagnosis	Chest X-ray images	CNN, ResNet-18.[18]
Classification	Diabetic Retinopathy Diagnosis	Fundus images	CNN, MobileNet-v2.[19]
Segmentation	Brain Tumor Segmentation	MRI scans	CNN, U-Net.[20]
Segmentation	Cardiac MRI Segmentation	MRI scans	CNN, U-Net.[21]
Segmentation	Lung Tumor Segmentation	CT scans	CNN, U-Net.[22]
Regression	Mortality Prediction	Electronic Health Records	LSTM, MLP [23]
Regression	Blood Pressure Prediction	PPG signals	LSTM.[24]
Regression	Heart Failure Prediction	Electronic Health Records	Multi-task LSTM .[25]

SHEFL models, for example, scored 80.32% on the Dice, which is quite similar to the 81.71% of FL models. However, training may be slowed down by homomorphic encryption's computational overhead. Furthermore, SHEFL enhances privacy but does not completely remove risks. Collaboration between institutions and a secure infrastructure are also prerequisites for the method to work.

Models like M-VGG16, M-ResNet50, M-ResNet101, and ViT were proposed by Ahsan et al.

[26] to classify monkeypox cases using image analysis. On a small dataset, the M-VGG16 achieved 88% accuracy, and on an imbalanced dataset, it achieved 76-77%. With 89% accuracy, the M-ResNet50 achieved the best results on a multiclass dataset. Adam-trained models outperformed SGD-trained models. The research emphasizes how FL can facilitate cooperative model training without data sharing, protecting patient privacy, and enhancing AI-based diagnostic models. The potential for creating reliable and secure healthcare solutions is demonstrated by the integration of FL. The study is subject to various limitations, such as the scarcity of images depicting monkeypox, the requirement for expert verification, testing on data that is significantly unbalanced, and the lack of mobile diagnostic tools.

V

Challenges of Federated Learning

Although FL provides the ability for multiple devices to work together on ML model training without exchanging raw data, there are a number of issues that need to be resolved. These noteworthy difficulties consist of:

Heterogeneity of Data and Devices:

It is difficult to deploy FL in the healthcare sector effectively due to the variety of data and devices.

- **Data heterogeneity:** There are many different types of healthcare data, including text, images, and time series. Each type of data has its own properties that call for particular processing techniques in order to train models. Difficulties arise from differences in data quality between devices, which can be attributed to patient demographics, sensor accuracy, and device specifications. For FL to be successful, standardization and data pre-processing are therefore crucial.
- **Device Heterogeneity:** The power consumption, network connectivity, and hardware and software specifications of healthcare devices differ. Certain devices might not have the computational or memory capacity to take part in FL due to resource constraints. Additional data interoperability and communication efforts are also needed due to differences in operating systems or programming languages between devices. To address this heterogeneity in FL and enable healthcare to overcome variations in data and devices, a number of strategies have been developed.
- **Federated Transfer Learning:** To expedite model training on heterogeneous devices, this technique leverages pre-trained models on related data domains. The model parameters can be changed by adaptive learning algorithms based on the processing power of the device. It is feasible to reduce device computing load and data transfer by utilizing effective communication protocols.

Data privacy and security:

While there are many advantages to this approach, such as increased privacy and reduced communication costs, there are also serious issues with data protection. A few of the major problems with data privacy that FL raises are as follows:

- **Data leakage:** In FL, devices distribute updates to the model, but these changes could still include private information about the local data. Attackers might be able to intercept these updates. To prevent data leaks, privacy-preserving techniques like encryption must be applied.
- **Model inversion attacks:** FL is susceptible to this privacy issue as well. These attacks replicate the initial training data used by the local devices using the model updates. To prevent these attacks, privacy-preserving protocols that protect local data must be implemented.
- **Membership inference attacks:** If an attacker has access to specific devices, they might be able to determine which device was used to train the model in FL. This information could potentially reveal sensitive information about the user of the device or local data. Membership inference attacks can be prevented by using privacy-preserving techniques such as differential privacy.
- **Attacks on the central server:** In FL, the central server gathers model updates from the local devices. If the primary server is compromised, an attacker could obtain access to all updates and infer important details about the local data. For this reason, strong security measures must be implemented in order to protect the central server.

In general, maintaining data privacy presents a big challenge in Florida. A combination of privacy-preserving techniques, robust security measures, and meticulous FL system design and implementation are needed to overcome these obstacles. Furthermore, Figure 6 illustrates FL's difficulties in various health-related fields.

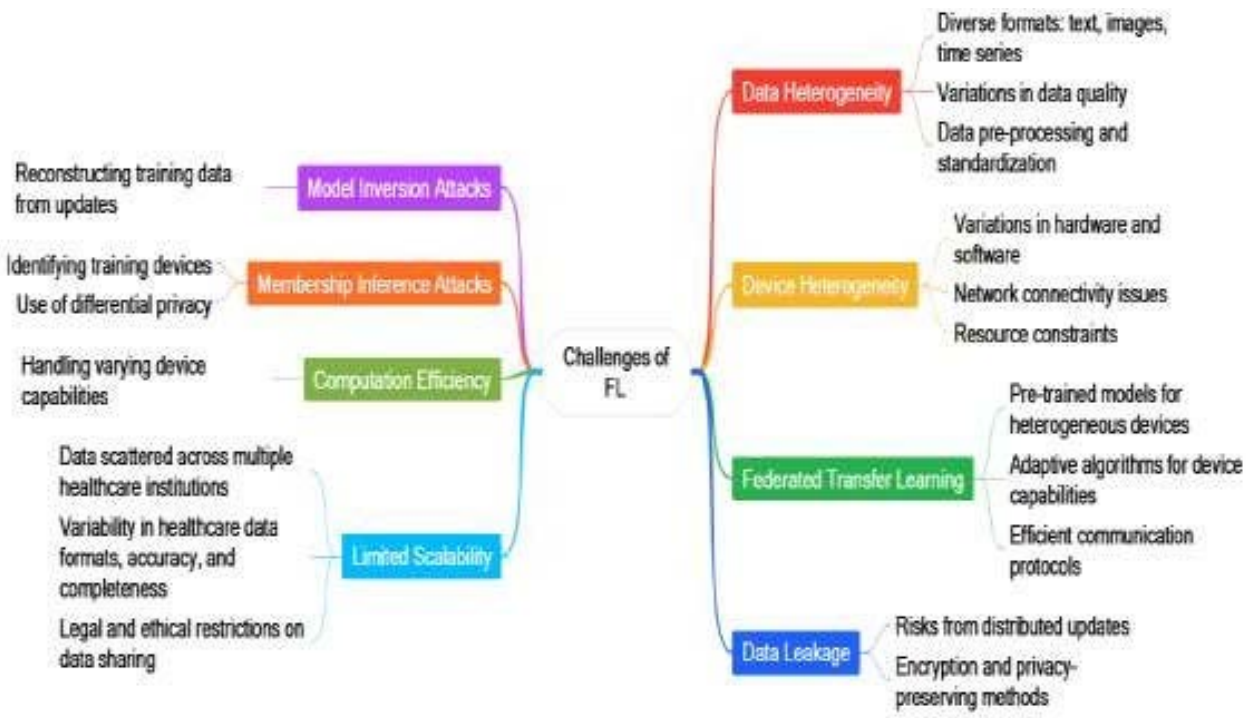

Figure 6: Challenges of FL in healthcare.

Communication and computation efficiency

Since the data from the participating devices needs to be sent securely and quickly to the central server for model training, communication effectiveness is a critical issue for FL. In order to ensure that the models are updated as quickly as possible and to protect user privacy, the data must be delivered with the least amount of delay possible. But FL also requires assistance with computing effectiveness. Therefore, depending on how powerful the participating devices are, the model's performance may change. The limited resources of the device must also be used to train the models, which could lead to excessive battery usage. FL techniques are updated frequently to fix these problems. One solution to the communication and computation efficiency issue is to use model compression techniques to reduce the size of the model and potentially speed up communication.

To calculate efficiency, one can also use model parallelism, data partitioning, and selective participation. In summary, there are significant communication and computation efficiency challenges with FL that need to be resolved in order to improve the technique's overall efficacy. To maintain FL as a competitive option for ML, researchers and practitioners are always seeking methods to enhance communication and computational efficiency.

Handling Non-IID (Independent and Identically Distributed) data:

One of FL's biggest problems is non-IID (Independent and Identically Distributed) data handling. In machine learning, the training data has traditionally been taken to be independent and uniformly distributed (IID). With FL, on the other hand, data is collected from multiple sources, some of which must be more impartial or dispersed.

The non-IID nature of the data creates a number of challenges in FL. For example, data distribution can differ dramatically amongst customers, which makes it challenging to train a general model that works well for every customer. Non-IID data may introduce bias into training, which could result in subpar output.

Various techniques such as data augmentation, transfer learning, and client weighing have been proposed to address these challenges. The process of client weighting entails dividing up a client's weights according to how their data is distributed. By using this strategy, the training process can concentrate more on clients with more representative data, which enhances the overall performance of the model. In order to help make the data more representative, data augmentation involves generating new data samples and adding noise or making small adjustments to the existing data.

A few studies that have tried to develop effective FL algorithms for non-IID data are FedProx, SCAFFOLD, and FedNova. One transfer learning technique that can assist in overcoming the difficulties presented by non-IID data is to use a pre-trained model as a foundation for training on the non-IID data.

In conclusion, handling non-IID data is one of FL's primary challenges. Several approaches, including client weighting, data augmentation, and transfer learning, have been used to tackle this issue. To develop more useful strategies for handling non-IID data in FL, more research is required.

Data Bias

Data bias is the term used to describe when the model's training set contains inaccurate or biased data. This might occur if certain servers or devices have different types of data than others, if the data isn't sampled randomly, or if some devices or servers have more data than others. This could lead to the model becoming biased in favour of the training set, which would negatively affect how well it performed on new, untrained

data. In Florida, a variety of factors can cause data bias. For instance, the model might perform poorly on devices with

different allocations and become biased towards devices with a different data distribution or more data than others.

Furthermore, if the data is skewed or not randomly sampled, the final model might not accurately represent the entire population. There are several approaches to address the problem of data bias in FL. One procedure is to make sure the data's sample size and population representation are random. An alternative strategy is to use data augmentation techniques to generate additional training data that is representative of the entire population. Additionally, devices with different or less data can be given a higher weight in FL algorithms, or the model can be adjusted using adaptive learning algorithms that modify the model based on the distribution of data on each device. These methods allow FL algorithms to account for data bias.

Limited scalability

The limited scalability of FL in the healthcare industry is one of its problems. The distribution of healthcare data among various hospitals, clinics, and other healthcare providers often poses a challenge in terms of organizing the data aggregation and sharing needed for FL. An additional challenge is the erratic nature of medical data. The format, completeness, and accuracy of healthcare data vary, which could affect how FL-trained ML models function. The amount of healthcare data that can be shared and used to train machine learning models may be restricted by ethical and legal constraints. Researchers are investigating various strategies to increase the scalability of FL in healthcare to tackle these issues. One of these is developing policies and processes for provider cooperation and information sharing. The creation of privacy-preserving protocols to safeguard sensitive data, the development of innovative algorithms to manage distributed and heterogeneous data, and some other examples. Stakeholders like legislators, patients, and providers must be involved in order to guarantee that the risks and benefits of FL are fairly balanced and to promote confidence in its use in the healthcare sector.

Lack of interpretability

The capacity to understand how a model arrives at its conclusions or forecasts is known as interpretability. It's crucial for a number of reasons, including accountability, transparency, and trust. However, because FL uses a centralized server to aggregate model updates from distributed raw data, it is challenging to understand the entire model. The interpretability issue in FL can be fixed by employing techniques like differential privacy, which introduces noise into the data to safeguard individual privacy while still allowing the model to be appropriately trained. Using techniques like model distillation is an alternate tactic. Using the same set of data, a more basic model serves as the foundation for training the global model. FL generally presents interpretability challenges, but there are a number of tactics that can be used to get around these problems and make the creation easier.

VI

Possible Future Research Directions

FL can be fully utilized to improve healthcare outcomes while protecting patient privacy if these research avenues are pursued. There are a plethora of potential directions for FL research in the medical domain, including:

- **Developing strategies for privacy protection:** Because patient data is sensitive and needs to be secure, future research in FL in healthcare will have to prioritize developing privacy- preserving techniques.
- **Model generalization:** ML models developed using FL may be more likely to overfit, which may limit their usefulness in real-world situations. Improving FL model generalization can ensure the accuracy and efficacy of these models when applied in healthcare settings. To optimize FL models' performance in real-world healthcare settings, researchers need to focus on improving model generalization and reducing over fitting.
- **Identifying and addressing data biases:** Biases in healthcare data have the potential to worsen and prolong existing disparities and inequities. These models have the potential to reduce healthcare disparities by identifying and resolving data bias in FL. Thus, addressing data biases may be a primary focus of future FL research in the healthcare industry.
- **Creating novel architectures for tasks unique to the healthcare industry:** The distinct data and task requirements found in the healthcare sector may require special FL model architectures. We need to develop new models and architectures for better performance.
- **Integrating FL with other technologies:** By integrating FL with other technologies

like edge computing and block chains, new healthcare solutions can be created. Future technologies can assist in addressing specific unique opportunities and challenges in the healthcare sector when combined with FL.

- **Using FL in low-resource environments:** Where data and computer resources may be limited, FL is an effective tool. Future FL research in the healthcare sector may find it to be an important topic given its potential to improve healthcare outcomes in underprivileged areas.

Overall, there are many exciting opportunities for additional FL research in the medical domain. As the technology advances, researchers will need to address these problems and develop new techniques and resources in order to enable the widespread use of FL in healthcare.

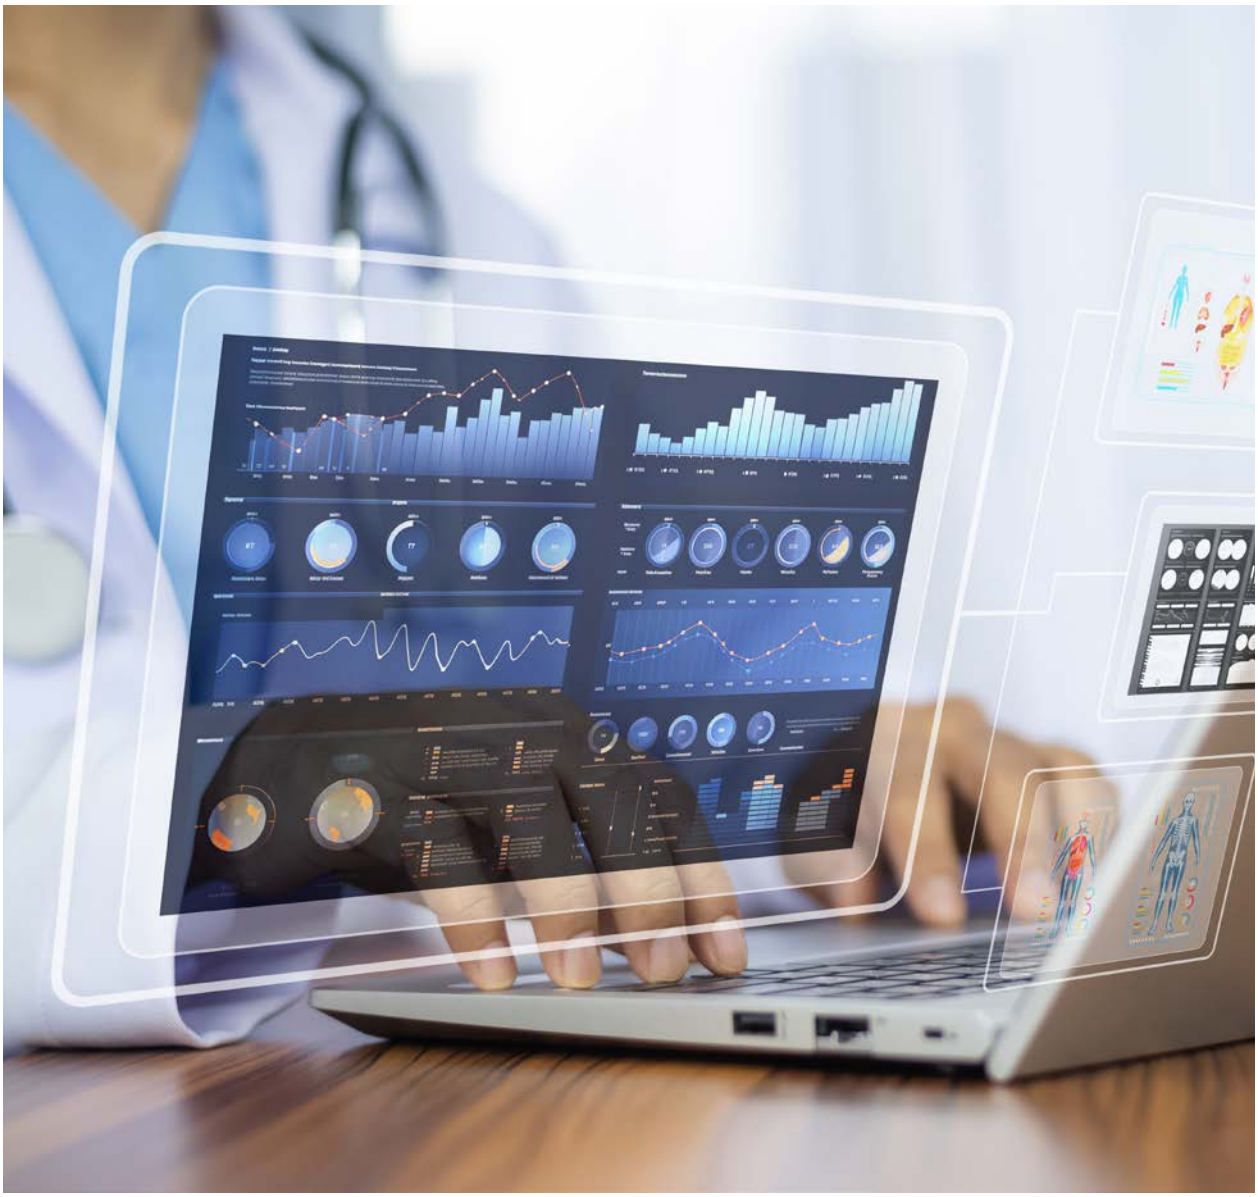

VII. Conclusion

In this paper, we investigate the use of Federated Learning (FL) in the healthcare industry, emphasizing how it can improve security and privacy when managing private medical data. Our results demonstrate that FL, by utilizing techniques like differential privacy and secure aggregation, can effectively mitigate risks like data leakage and model inversion attacks. There are still issues, though, mainly with handling heterogeneous data and devices, making sure that communication and computation are efficient, and dealing with the non-IID character of healthcare data. These difficulties make FL deployment more difficult and provide major logistical, ethical, and technical obstacles to its wider adoption. Thus, future research should concentrate on creating scalable, reliable algorithms that can effectively manage distributed, heterogeneous datasets while preserving patient confidentiality and data security. It is imperative that future research concentrate on the interpretability and fairness of FL models, guaranteeing their accessibility and equity for a wide range of patient demographics. By resolving these problems, FL will operate more efficiently and gain the trust of its stakeholders, which will promote its integration into standard medical procedures. Successful FL integration has the potential to have a significant impact on healthcare, revolutionizing patient care by enabling more preventive and personalized medicine while protecting patient privacy.

References

[1] Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. Federated machine learning: Concept and applications. ACM Transactions on Intelligent Systems and Technology (TIST), 10(2):1–19, 2019.

[2] Muhammad Mateen Yaqoob, Muhammad Nazir, Muhammad Amir Khan, Sajida Qureshi, and Amal Al-Rasheed. Hybrid classifier-based federated learning in health service providers for cardiovascular disease prediction. Applied Sciences, 13(3):1911, 2023.

[3] Akhil Vaid, Suraj K Jaladanki, Jie Xu, Shelly Teng, Arvind Kumar, Samuel Lee, Sulaiman Somani, Ishan Paranjpe, Jessica K De Freitas, Tingyi Wanyan, et al. Federated learning of electronic health records to improve mortality prediction in hospitalized patients with covid-19: machine learning approach. JMIR medical informatics, 9(1):e24207, 2021.

[4] Gokberk Elmas, Salman UH Dar, Yilmaz Korkmaz, Emir Ceyani, Burak Susam, Muzaffer Ozbey, Salman Avestimehr, and Tolga Çukur. Federated learning of generative image priors for mri reconstruction. IEEE Transactions on Medical Imaging, 2022.

[5] Rajesh Kumar, Abdullah Aman Khan, Jay Kumar, Noorbakhsh Amiri Golilarz, Simin Zhang, Yang Ting, Chengyu Zheng, Wenyong Wang, et al. Blockchain- federated-learning and deep learning models for covid-19 detection using ct imaging. IEEE Sensors Journal, 21(14):16301–16314, 2021.

[6] Xiaoxiao Li, Yufeng Gu, Nicha Dvornek, Lawrence H Staib, Pamela Ventola, and James S Duncan. Multi-site fmri analysis using privacy-preserving federated learning and domain adaptation: Abide results. Medical Image Analysis, 65:101765, 2020.

[7] Weishan Zhang, Tao Zhou, Qinghua Lu, Xiao Wang, Chunsheng Zhu, Haoyun Sun, Zhipeng Wang, Sin Kit Lo, and Fei-Yue Wang. Dynamic-fusion-based federated learning for covid-19 detection. IEEE Internet of Things Journal, 8(21):15884–15891, 2021.

[8] Bless Lord Y Agbley, Jianping Li, Amin Ul Haq, Edem Kwedzo Bankas, Sultan Ahmad, Isaac Osei Agyemang, Delanyo Kulevome, Waldiodio David Ndiaye, Bernard Cobbinah, and Shoistamo Latipova. Multimodal melanoma detection with federated learning. In 2021 18th International Computer Conference on Wavelet Active Media Technology and Information Processing (ICCWAMTIP), pages 238–244. IEEE, 2021.

[9] Adnan Qayyum, Muhammad Umar Janjua, and Junaid Qadir. Making federated learning robust to adversarial attacks by learning data and model association. Computers & Security, 121:102827, 2022.

[10] Shinu M Rajagopal, M Supriya, and Rajkumar Buyya. Fedsdm: Federated learning based smart decision making module for ecg data in iot integrated edge-fog-cloud computing environments. Internet of Things, page 100784, 2023.

[11] Satyabrata Aich, Nday Kabulo Sinai, Saurabh Kumar, Mohammed Ali, Yu Ran Choi, Moon-IL Joo, and Hee-Cheol Kim. Protecting personal healthcare record using blockchain & federated learning technologies. In 2022 24th International Conference on Advanced Communication Technology (ICACT), pages 109–112. IEEE, 2022.

[12] Yaojie Wang, Xiaolong Cui, Zhiqiang Gao, and Bo Gan. Fed-scnn: a federated shallow-cnn recognition framework for distracted driving. Security and Communication Networks, 2020:1–10, 2020.

[13] Madhura Joshi, Ankit Pal, and Malaikannan Sankarasubbu. Federated learning for healthcare domain-pipeline, applications and challenges. ACM Transactions on Computing for Healthcare, 3(4):1–36, 2022.

[14] Ahmet Ali Süzen and Mehmet Ali ̧Sim ̧sek. A novel approach to machine learning application to protection privacy data in healthcare: Federated learning. Namık Kemal Tıp Dergisi, 2020.

[15] Qiong Wu, Kaiwen He, and Xu Chen. Personalized federated learning for intelligent iot applications: A cloud-edge based framework. IEEE Open Journal of the Computer Society, 1:35–44, 2020

[16] Nazia Shehnaz Joynab, Muhammad Nazrul Islam, Ramiza Rumaiza Aliya, ASM Rakibul Hasan, Nafiz Imtiaz Khan, and Iqbal H Sarker. A federated learning aided system for classifying cervical cancer using pap-smear images. Informatics in Medicine Unlocked, page 101496, 2024.

[17] Daniel Truhn, Soroosh Tayebi Arasteh, Oliver Lester Saldanha, Gustav Müller-Franzes, Firas Khader, Philip Quirke, Nicholas P West, Richard Gray, Gordon GA Hutchins, Jacqueline A James, et al. Encrypted federated learning for secure decentralized collaboration in cancer image analysis. Medical image analysis, 92:103059, 2024.

[18] Joaquim de Moura, Jorge Novo, and Marcos Ortega. Fully automatic deep convolutional approaches for the analysis of covid-19 using chest x-ray images. Applied Soft Computing, 115:108190, 2022.

[19] Zhen Ling Teo, Aaron Y Lee, Peter Campbell, RV Paul Chan, and Daniel SW Ting. Developments in artificial intelligence for ophthalmology: Federated learning, 2022.

[20] Patrick Foley, Micah J Sheller, Brandon Edwards, Sarthak Pati, Walter Riviera, Mansi Sharma, Prakash Narayana Moorthy, Shih-han Wang, Jason Martin, Parsa Mirhaji, et al. Openfl: the open federated learning library. Physics in Medicine & Biology, 67(21):214001, 2022.

[21] Alan C Kwan, Gerran Salto, Susan Cheng, and David Ouyang. Artificial intelligence in computer vision: cardiac mri and multimodality imaging segmentation. Current cardiovascular risk reports, 15:1–8, 2021.

[22] Sajid Nazir and Mohammad Kaleem. Federated learning for medical image analysis with deep neural networks. Diagnostics, 13(9):1532, 2023.

[23] uraj Rajendran, Zhenxing Xu, Weishen Pan, Arnab Ghosh, and Fei Wang. Data heterogeneity in federated learning with electronic health records: Case studies of risk prediction for acute kidney injury and sepsis diseases in critical care. PLOS Digital Health, 2(3):e0000117, 2023.

[24] Eoin Brophy, Maarten De Vos, Geraldine Boylan, and Tomas Ward. Estimation of continuous blood pressure from ppg via a federated learning approach. Sensors, 21(18):6311, 2021.

[25] Sawsan AbdulRahman, Hanine Tout, Hakima Ould-Slimane, Azzam Mourad, Chamseddine Talhi, and Mohsen Guizani. A survey on federated learning: The journey from centralized to distributed on-site learning and beyond. IEEE Internet of Things Journal, 8(7):5476–5497, 2020.

[26] Md Manjurul Ahsan, Tasfiq E Alam, Mohd Ariful Haque, Md Shahin Ali, Rakib Hossain Rifat, Abdullah Al Nomaan Nafi, Md Maruf Hossain, and Md Khairul Islam. Enhancing monkeypox diagnosis and explanation through modified transfer learning, vision transformers, and federated learning. Informatics in Medicine Unlocked, 45:101449, 2024.

National Centre of Excellence
CYBERSECURITY TECHNOLOGY AND ENTREPRENEURSHIP

The National Centre of Excellence (NCoE) for Cybersecurity Technology Development is a joint initiative between the Ministry of Electronics & Information Technology (MeitY), Government of India, and the Data Security Council of India (DSCI). Its primary objective is to catalyze and accelerate cybersecurity technology development and entrepreneurship within the country. NCoE plays a crucial role in scaling up and advancing the cybersecurity ecosystem, focusing on various critical and emerging security areas. Equipped with stateof-the-art facilities, including advanced lab infrastructure and test beds, NCoE facilitates research, technology development, and solution validation for adoption across government and industrial sectors. Adopting a concerted strategy, NCoE endeavors to translate innovations and research into market-ready deployable solutions, thereby contributing to the evolution of an integrated technology stack comprising cutting-edge, homegrown security products and solutions.

DSCI
PROMOTING DATA PROTECTION
A nasscom Initiative

Data Security Council of India (DSCI) is a premier industry body on data protection in India, setup by nasscom, committed to making the cyberspace safe, secure and trusted by establishing best practices, standards and initiatives in cybersecurity and privacy. DSCI brings together governments and their agencies, industry sectors including ITBPM, BFSI, telecom, industry associations, data protection authorities and think-tanks for policy advocacy, thought leadership, capacity building and outreach initiatives. For more info, please visit www.dsci.in

DATA SECURITY COUNCIL OF INDIA

+91-120-4990253 | ncoe@dsci.in

https://www.n-coe.in/

4 Floor, NASSCOM Campus, Plot No. 7-10, Sector 126, Noida, UP -201303

Follow us on

@CoeNational
nationalcoe
nationalcoe
NationalCoE