

A Communication-Efficient Federated Learning Framework FOR Privacy-Preserving Disease Diagnosis IN Low-Resource Healthcare Systems

Vinodhini Ravikumar¹

Submitted: 02/04/2026

Revised: 14/05/2026

Accepted: 25/05/2026

Abstract: However, in low-resource countries, implementing A.I.-driven diagnostic systems is a daunting task because of privacy issues, restrictions in computational resources and network reliability. Current centralized machine learning approaches involve sharing of sensitive patient information with external servers, which raises privacy concerns and restricts collaborative healthcare analytics. In this paper, the authors present a communication efficient federated learning model for privacy-preserving disease detection in resource-limited healthcare systems. It combines lightweight convolutional neural networks with secure parameter aggregation to minimize the amount of communication needed without compromising diagnostic performance. We have performed experiments with the Chest X-ray pneumonia dataset distributed to simulated healthcare clients. The suggested model was found to be accurate in diagnostics with a value of 94.1% and the communication cost was reduced by 36% from the traditional federated learning model. The study presents a privacy-preserving and scalable framework for health care AI that is appropriate for decentralized medical diagnosis in remote areas.

Keywords: Federated Learning; Privacy-Preserving AI; Medical Diagnosis; Healthcare Informatics; Low-Resource Healthcare; Artificial Intelligence in Healthcare

1. Introduction

1.1 Background

Healthcare Industry is seeing a rapid transformation with the use of Artificial Intelligence (AI) which is allowing sophisticated data analysis, intelligent clinical decisions making and automated disease diagnosis. The advent of machine learning, deep learning, and predictive analytics has revolutionized the Healthcare sector, its delivery, diagnostic capabilities, and patient management systems (Secinaro et al., 2021; Davenport & Kalakota, 2019). The application of AI technologies in medical imaging and diagnosis, in the field of pathology to facilitate precise tissue analysis, in nursing care for enhanced patient support, in patient monitoring to streamline patient care processes, and in personalized medicine to guide tailored treatments is growing and shows potential to significantly improve clinical efficiency and healthcare outcomes in radiology and other areas of the healthcare sector (Banerjee & Kumar, 2024; Koo et al., 2024). Cloud computing systems, medical image databases, and electronic health records have also become increasingly commonplace, propelling AI's use in healthcare. Furthermore, the availability of electronic health records, medical imaging databases, and cloud computing infrastructures has further fueled the application of AI in healthcare (Alzghoul, 2024).

One of the most impactful applications of AI in the healthcare industry is medical diagnosis automation. The deep learning models, especially convolutional neural networks

(CNNs), have shown outstanding result in the medical image processing and clinical data analysis to identify disease like pneumonia, skin cancer, cardiovascular abnormality, diabetic complication etc. (Wen & Huang, 2022; Alhejaily, 2025). In high-pressured healthcare environments, automated diagnostic systems can help healthcare professionals to cut down on diagnostic delays, enhance the accuracy of diagnosis and facilitate medical processes. Automated diagnostic systems can help professionals in healthcare to achieve the following things: reduce diagnostic delays, improve diagnostic accuracy, and support medical processes, particularly in high-pressure healthcare environments (Dave & Patel, 2023). Moreover, AI-powered health care systems can enhance health care accessibility in communities with limited access to healthcare services and lack of medical experts (Francisca Chibugo Udegbe et al., 2024).

While AI systems in healthcare have significant benefits, they are also accompanied by various technical, ethical, and operational challenges. The explainability, fairness, trustworthiness, and patient privacy concerns are still hindering the general use of AI technologies in healthcare (Amann et al., 2020; Sadeghi et al., 2024). Many healthcare professionals and patients are concerned with the sensitivity of the data collected, processed, and shared in centralized AI systems (Fritsch et al., 2022; Hatem et al., 2024). Additionally, ethical issues related to using AI in healthcare, such as transparency, accountability, and data management, have been gaining traction in the scholarly community (Prakash et al., 2022; Vandemeulebroucke, 2025).

Federated learning (FL) has been proposed as an alternative distributed learning method to solve these issues, which involves training a joint model with participation from multiple data centers while keeping patient data from each center confidential

Founder at Mind Mosaic AI, Vancouver, - V5K3K2, BC, Canada
Email: vini@mindmosaicai.com

(Xu et al., 2021). In federated learning, healthcare institutions will train ML models with local data, and only share model parameters or updates to a central aggregating server (Zhang et al., 2021). This decentralized data collection approach drastically minimizes the risks of centralized data gathering and enables privacy protecting healthcare analytics (Dhade & Shirke, 2023). In healthcare informatics, federated learning is particularly promising because of its ability to facilitate secure, scalable, and collaborative AI systems in healthcare, which aim to improve patient outcomes and provide value to healthcare systems and society as a whole. In the healthcare informatics field, federated learning holds significant potential for enabling secure, scalable, and collaborative healthcare AI systems that can enhance patient outcomes and deliver value to healthcare systems and society.

Additional research work in recent times has highlighted the increasing adoption of federated learning in various medical image applications, including disease prediction, mental health, and decentralized healthcare (Upreti et al., 2024; Vajrobol et al., 2025). The use of FL in healthcare settings has been significantly enhanced by innovations in federated optimization, secure aggregation, and communication-efficient training techniques (Ji et al., 2024; Zhang et al., 2024). Furthermore, to improve the security of federated healthcare systems, research has been conducted on privacy-preserving approaches like differential privacy and aggregation algorithms based on encryption (Hu et al., 2024; Li et al., 2023).

1.2 Problem Statement

While centralized AI systems show promise in diagnostic capabilities, their adoption in healthcare settings faces significant challenges, including data privacy, regulatory compliance, and infrastructure needs. Conventional machine learning applications mandate that hospitals and healthcare facilities share patient data with a central cloud system for model training and analysis, which poses risks of data breaches. Traditional machine learning platforms depend on hospitals or healthcare institutions uploading patient information to a central cloud system for modelling and analysis, risking data breaches. These centralized architectures can pose significant challenges to data security, privacy, and access, especially for sensitive medical data (Li et al., 2023). The interdisciplinary sharing of healthcare data is further complicated by regulatory frameworks and ethical guidelines concerning data protection of patients (Lekadir et al., 2025).

These challenges are even greater when implementing such measures in low resource healthcare organisations, such as in rural and underserved areas where healthcare provision is still poor. The implementation of extensive AI solutions is hindered in many healthcare facilities in developing areas due to limited technical skills and knowledge, as well as inadequate high-speed internet connectivity and computing power (Aljehani & Al Naweess, 2025). Distributed healthcare AI systems can be impacted by factors such as limited bandwidth, unstable network environments, and insufficient computational resources, which can hinder their efficiency and reliability. For this reason, numerous current AI health care solutions cannot be used by facilities with limited resources.

Moreover, the current federated learning systems generally resort to ideal communication and computation environments which may not be the case in resource-limited healthcare environments. However, most federated learning frameworks use a number of communication rounds, heavy model updates

and deep learning architectures with high computing complexity that consume more energy and cause more communication overheads (Liu et al., 2024). This approach is not feasible in scenarios with limited connectivity and hardware resources, such as in remote healthcare settings. In this context, it is important to have lightweight, communication-efficient and privacy-preserving federated learning frameworks tailored for low-resource healthcare settings.

1.3 Research Gap

Prior studies in federated learning for healthcare have predominantly centered on enhancing diagnostic performance, data privacy and security, and federated collaboration among healthcare institutions (Xu et al., 2021; Wen et al., 2023). These research efforts have helped to pave the way for privacy-preserving healthcare AI, but there are several key gaps that need to be addressed.

The first is that most federated learning models have the premise of stable and high-speed internet connections between the participating healthcare institutions. This is not a feasible scenario for rural healthcare institutions or for implementing a healthcare system in transitional regions in which the communication network is not always reliable and bandwidth is scarce (Upreti et al. 2024). In a traditional federated learning setting, it is therefore possible that there is too much communication overhead and training latency as a result of frequent parameter transmission.

Second, some federated learning methods use complex deep learning models which demand high memory and computational requirements, necessitating high-performance hardware accelerators (Ji et al., 2024). These high resource consumption architectures are hard to adopt in lower resource healthcare settings where hospitals can have restricted computing resources.

Third, previous research focuses mainly on the privacy protection and model performance, neglecting the real-world deployment issues in under-resourced healthcare environments (Li et al., 2025). Some challenges like effective communication, light-weight model optimization, and FL using resource-constrained resources still need to be explored in the healthcare context.

Therefore, it is essential to develop a communication-efficient and lightweight FL framework that works in low-resource healthcare systems to facilitate privacy-preserving disease diagnosis while maintaining diagnostic accuracy.

1.4 Research Objective

In this study, we propose to create a low-computation and communication workload Federated Learning system for privacy-preserving disease diagnosis service in resource-constrained healthcare systems. The proposed framework aims to reduce the communication overhead, computational complexity and guarantee patient privacy while ensuring high diagnostic accuracy.

In particular, the project aims to:

- Create a distributed and federated learning system in the healthcare diagnosis.
- Optimize processes to minimize network overheads by using communication efficient methods.

- Maintain privacy of patient's data with secure mechanisms of parameter aggregation.
- Test the proposed framework on a real world medical imaging dataset.
- Compare the proposed model with conventional, centralized and federated model.

1.5 Contributions

This study has a number of significant contributions, which are summarized:

- Proposal of the lightweight federated learning architecture of decentralized healthcare diagnosis.
- An optimization method is embedded to optimize the communication resources, in the case of low bandwidth in healthcare applications.
- Data aggregation mechanisms are included that preserve privacy to improve the security of patient data.
- A medical image dataset of chest X-rays is used to evaluate the proposed framework.
- Comparative experiments show enhanced communication efficiency, and the high diagnostic accuracy.
- The research offers real-world lessons learned on implementing federated healthcare AI systems, in particular, within low-resource and underserved healthcare environments.

2. Related Work

2.1 AI-Based Disease Diagnosis

The impact of AI on healthcare is particularly profound in the field of disease diagnosis, where managing vast amounts of medical information is both faster, more accurate, and efficient than ever before. As machine learning and deep learning technologies have become more prevalent in the field of healthcare, they have been used to improve decision-making in clinical practice, interpretation of medical images, and predictive analytics (Davenport and Kalakota, 2019; Secinaro et al., 2021). In radiology, pathology, oncology, cardiology, and infectious disease detection, AI systems are increasingly being adopted in the medical field, enhancing healthcare quality and minimizing diagnostic delays (Alzghoul, 2024).

Deep learning is one of the most promising AI methods for the automation of medical diagnosis because it can learn intricate hierarchical representations from medical data. CNNs have proven to be highly effective for image-based disease diagnosis tasks, especially in medical imaging, like chest X-ray analysis, skin lesion classification, brain tumor detection, and retinal disease screening (Wen & Huang, 2022). CNNs automatically learn significant spatial features in medical images, which alleviates the reliance on manual feature engineering and enhances classification performance.

In recent years, with the advent of sophisticated architectures in deep learning, the applications of AI in healthcare have been further enhanced with the ability to process heterogeneous healthcare data (Banerjee & Kumar, 2024). CNN-based systems have been effective in identifying

pneumonia from chest X-ray images and abnormalities from computed tomography and magnetic resonance images. These models can help in early detection of the disease and aid in large scale diagnostic screening programs, especially in areas where medical experts are scarce (Dave & Patel, 2023).

There are also several studies that have focused on the wider implications of AI systems in healthcare on healthcare delivery and efficiency. Balpande et al. (2025) noted that AI is still transforming healthcare and nursing practices by providing intelligent automation and predictive healthcare analytics. Likewise, Chen (2025) highlighted the increasing role of AI in doctor-patient interactions and healthcare administration in urban settings in Asia. A detailed bibliometric analysis by Jimma (2023) and Xie et al. (2024) also showed that within the past 10 years, the research on AI in healthcare has rapidly increased.

However, there are still many challenges to AI-driven healthcare systems in terms of transparency, fairness, explainability, and trustworthiness. One of the problems with deep learning models is that their decision-making process is opaque for health care professionals, which is why they are often referred to as "black-box" systems (Amann et al., 2020). In the context of healthcare, the importance of explainability in AI systems is increasingly recognised, especially in clinical settings where accountability and interpretability are crucial for patient safety and regulatory adherence (Sadeghi et al., 2024). Furthermore, the issue of fairness and bias in AI-based healthcare systems has come to the fore, with the potential for unequal access to healthcare due to biased models (Ueda et al., 2024).

Beyond technical barriers, healthcare organizations are also hesitant to implement AI systems because of ethical, legal, and privacy issues. There are often concerns in the healthcare sector regarding the privacy and security of the collection, processing, and use of sensitive medical data by AI systems, as well as their potential impact on patient privacy and the ability to maintain confidentiality (Hameed et al., 2023; Fritsch et al., 2022). Some ethical issues of AI in healthcare are data misuse, lack of transparency, algorithmic bias, and accountability (Prakash et al., 2022). AI-powered disease diagnosis provides significant advantages, but the development of secure, transparent, and privacy-preserving healthcare AI systems is a key research focus.

2.2 Federated Learning in Healthcare

With its distributed nature, federated learning provides a viable alternative to centralized healthcare AI solutions that can tackle multiple privacy and data ownership issues. In contrast to typical machine learning structures that demand the gathering of patient information in a centralized structure, federated learning permits healthcare organizations to collectively prepare machine learning models and not share sensitive medical records. Federated learning involves healthcare organizations cooperating to train machine learning models without sharing patient-specific data directly, thereby avoiding the centralization of medical records that might be sensitive. This decentralized method involves training AI models locally at healthcare facilities with their own data and only sending them updates of the model or its parameters to a central server for aggregation.

As patient data becomes more valuable, privacy and security concerns are rising, and the landscape of healthcare is becoming more complex due to the need for collaboration across multiple institutions, federated learning is becoming more

popular. Healthcare organizations may have data protection regulations which restrict external sharing of patient data. To tackle these challenges, federated learning has been introduced where hospitals or medical institutions can jointly gain from the distributed healthcare data, without relinquishing control of local data (Dhade & Shirke, 2023).

In the field of healthcare, recent research has shown the potential of FL for various applications such as medical image analysis, disease prediction, remote patient monitoring, and personalized medicine (Joshi et al., 2022). Chest X-ray (Upreti et al., 2024), cancer diagnosis, diabetic retinopathy detection, and mental healthcare analytics (Vajrobal et al., 2025) are a few examples that have successfully implemented federated learning frameworks. The applications demonstrate the capabilities of FL to enhance the accuracy of diagnosis, but at the same time maintain patient privacy.

Federated Learning (FL) has also garnered significant attention in healthcare informatics through several surveys. Given its capabilities of privacy preservation, Zhang et al. (2021) highlighted Federated Learning as one of the most important advancements in Distributed Artificial Intelligence. Wen et al. (2023) pointed out that federated learning enables the training of AI models for multiple institutions that are located in different geographical areas without the need to store data in a centralized location. Likewise, Liu et al. (2024) presented the state-of-the-art federated optimization, model aggregation, and distributed learning architectures that have been developed to facilitate scalability and communication efficiency.

Federated learning is one of the most beneficial applications in the healthcare environment, specifically collaborative hospital learning. With the cooperation and training, hospitals that do not have enough data can all participate in the formation of diagnostic models without infringing on patient privacy, which can enhance the accuracy of the diagnostic model (Ji et al., 2024). For rare diseases, it is often helpful to have access to a diverse range of data sources, or in low-resource healthcare settings where local data might not be sufficient to train highly accurate medical AI models.

Moreover, recent advances in the methodology of federated healthcare AI have added techniques like federated transfer learning, federated X learning and adaptive model aggregation to enhance the performance and interoperability of models across heterogeneous healthcare systems (Zhang et al., 2024). The advances showcase the increasing sophistication and viability of federated learning as a scalable approach to distributed healthcare intelligence.

Though it offers numerous benefits, federated learning continues to encounter a number of practical hurdles to be overcome in healthcare settings. Even today, issues such as communication overhead, client heterogeneity, computational complexity, and network environments that are unstable are still impacting the scalability and efficiency of federated healthcare systems (Li et al., 2025). These are still important challenges that must be overcome to allow federated learning to be used in real-world, resource-limited healthcare settings.

2.3 Privacy-Preserving AI Techniques

The need for privacy preservation is among the most important ones in Healthcare AI, thanks to the very sensitive nature of patient medical records. In addition to reducing the amount of direct information that is shared, model parameters and gradients

shared during collaborative training could contain confidential information unless adequate security measures are put in place (Hu et al., 2024). In response, a number of privacy-preserving methods have been created to enhance the security of federated healthcare systems.

Among the most commonly used privacy-preserving methods for federated learning, differential privacy is among them. Before they are sent to the aggregation server, differential privacy adds carefully calibrated noise to the model updates and this helps to minimize the risk of reconstructing the sensitive patient data from the shared parameters (Li et al., 2023). This way, the individual patient records can not be easily deduced from collaborative learning outputs, but the model performance is still acceptable.

Federated healthcare AI has also garnered a lot of interest in security aggregation mechanisms. Secure aggregation protocols can combine multiple updates to the model to create a unified model without exposing how each client contributed to the model update to the central server (Hu et al., 2024). Mechanisms like that enhance confidentiality and diminish the chances of being able to inspect the model during the distributed learning process without permission.

Federated healthcare environments can also benefit from the use of encryption-based techniques to preserve privacy. For distributed healthcare systems, the possibility of using homomorphic encryption for secure model training and secure parameter exchange has been investigated (Li et al., 2023). These techniques enable computation with encrypted data but prevent revealing information about the patients. Encryption-based techniques offer robust privacy protections, but can also require extra computation and communication.

The need for trustworthy and explainable healthcare AI systems that strike a balance between privacy protection and transparency and fairness is also a recent priority. Lekadir et al. (2025) introduced recommendations for deployable and reliable healthcare AI systems, emphasising the design principles of privacy, accountability, and interpretability. Likewise, Vandemeulebroucke (2025) highlighted the moral dimensions of AI in healthcare, and the requirement for AI governance frameworks in healthcare that are responsible at the international level.

Even with these advancements, there is still a great challenge in applying privacy preserving techniques in low resource healthcare settings, because of constraints on computing power and network infrastructure. For effective deployment of federated health care solutions in less developed areas, lightweight and communication efficient privacy preserving solution are thus required.

2.4 Issues in low-resource health care systems

Federated learning offers a new paradigm for privacy-preserving healthcare AI, but implementing federated learning in low-resource healthcare settings is challenging. In developing countries, many healthcare facilities suffer from various infrastructure constraints, such as limited internet access, computational power, technical knowledge, and funding for healthcare facilities (Aljehani & Al Nawees, 2025).

Bandwidth and network connectivity are the major problems. The traditional federated learning systems need regular interactions between participating clients and aggregation servers in the iterative training process. For

healthcare systems with inconsistent internet connectivity, frequent transmission of massive model parameters can greatly lead to high communication latency and low training efficiency (Liu et al., 2024). Interruption is a common problem that can also impact system reliability and model convergence.

However, another challenge is the availability of advanced computational tools like the GPU and high performance servers. However, many rural hospitals or small healthcare systems use low-power computing systems which are unable to effectively power computationally heavy deep learning models (Li et al., 2025). Thus, the deployment of large CNN architectures in the resource-limited healthcare setting could lead to significant processing delay and power consumption.

In addition, data access to limited labeled health care data is another major challenge in low-resource medical settings. Often the medical information at under-serviced health care facilities is not adequate and has not been annotated for high-performing AI models. Small local datasets could lead to overfitting the model and limit the generalization of the diagnostic to a broader patient population (Joshi et al., 2022). Optimizing and aggregating data in the federated model could be even more challenging due to the data heterogeneity found on various healthcare institutions.

Furthermore, there is a lack of technical skills in many low-resource healthcare environments to operate AI systems and distributed learning systems. Federated Healthcare Systems (FHS) are often hard to maintain in healthcare institutions because they lack suitably equipped AI engineers, cyber security professionals and data scientists to do so (Hameed et al., 2023). The challenges with these operations underscore the need for creating federated learning frameworks that are lightweight, scalable, and resource-efficient for use in real-life health care settings.

2.5 Research Gap Summary

Previous research has already shown the potential of artificial intelligence and federated learning in the areas of healthcare diagnosis, privacy protection and distributed collaboration in medicine. But, there are some issues yet to be addressed in existing healthcare federated learning studies. There are many well-established systems that are computationally complex, depend on stable communications, and focus on the accuracy of the models but do not take into account deployment issues in lower-resource healthcare. Moreover, the use of these privacy-preserving methods, like differential privacy and secure aggregation, can impose extra computational and communication cost which could restrict the feasibility in medical facilities with limited resources.

Thus, a lightweight and communication efficient federated healthcare framework that can facilitate privacy-sensitive disease diagnosis in limited infrastructure conditions is still in need. The current literature on this does not simultaneously consider privacy preservation, communication efficiency, and/or low resources deployment requirements.

3. Proposed Framework

This section introduces the proposed communication efficient Federated Learning (FL) framework for privacy-preserving disease diagnosis in low-resource healthcare settings. The framework is designed to combine decentralized training,

minimal deep learning, and privacy-preserving solutions for collaborative medical AI without sharing patient data.

3.1 System Architecture

The proposed system is based on the decentralized federated learning architecture which consists of multiple healthcare clients (hospitals/clinics) and a central aggregation server. The patient data in each healthcare institution is stored locally and diagnostic model is trained locally without transferring raw data from the institution.

These are the main components of the system:

- **Healthcare Clients (Hospitals/Clinics):**

Every client is a health center containing patient information from a specific area (such as chest x-rays, medical records). These clients perform local training using their private data.

- **Local Training Module:**

A light weight convolutional neural network (CNN) is fine-tuned at each institution using its own data. The model is trained on patterns that do not leak the sensitive information of patients.

- **Central Aggregation Server:**

The server orchestrates the FL process by collecting model updates (weights/gradients) from different clients and globally updating the model.

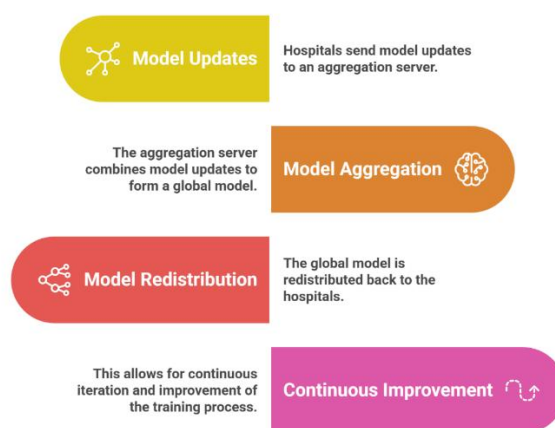
- **Secure Communication Channel:**

Provides secure communication of model parameters from client to server.

- **Global Model Distribution:**

The modified global model is sent to all participating clients for more training.

Conceptual Architecture Flow



This approach not only helps safeguard patient privacy by keeping raw medical data within the local healthcare institutions, but also allows for collaborative learning.

3.2 Federated Learning Workflow

The federated learning process in the proposed framework continues in an iterative manner and comprises of five primary steps:

(1) Local Model Initialization

A global model architecture is started at each of the participating hospitals. The parameters of the model are randomly initialized or pre-trained on a public medical dataset.

(2) Local Training Phase

The training of each client is done on their own private data set for a fixed number of epochs. In this stage, the model takes in the pattern of the hospital's patients in a localized pattern.

(3) Weight Transmission

Each client only reports new model parameters (weights/gradients) to the central server after local training. Not a single patient record is transferred.

(4) Global Aggregation

The central server fetches updates from the users and then applies a mechanism of weighted averaging (usually Federated Averaging, FedAvg). This yields a new global model which reflects knowledge from all the participating clients.

(5) Iterative Updating

This updated global model is communicated to all clients and the process repeats for several rounds of communication until convergence is reached.

This process is repeated and allows for distributed health systems to learn together while keeping data private.

3.3 Communication-Efficient Optimization

One drawback to the classic federated learning systems is the significant communication overhead, particularly when network connections are unstable or bandwidth is low. To overcome this, the framework proposed in the paper proposes a number of communication-efficient optimisation strategies. [23]

(1) Lightweight CNN Architecture

Instead of using deep learning architectures that are more complex and have more parameters, a compact convolutional neural network is used. This reduces:

- model size,
- local training time,
- Efficiently powering devices with limited resources.

The lightweight CNN is made feasible in rural hospitals, where GPUs or CPUs are limited.

(2) Gradient Sparsification

Only the most significant gradient changes are sent to the server to minimize the amount of traffic. Gradients are lowered where they are small or near zero, according to a set threshold. This significantly reduces:

- communication bandwidth usage,
- time for transmission per round.

(3) Model Compression Techniques

Model updates are compressed in the following ways:

- quantization (loss of accuracy of weights),
- Removing redundant parameters (pruning).

The techniques used shrink the amount of information sent without compromising on the accuracy of diagnosis.

(4) Reduced Communication Rounds

The framework defines more local training steps before communication, as opposed to a frequent synchronization. This reduces:

- The number of communication rounds,
- Reliance on regular internet service.

All of these optimizations create a system that can be used in low resource healthcare environments where network infrastructure is limited.

3.4 Privacy Preservation Mechanism

While federated learning naturally reduces the amount of raw data that is shared across multiple models, there are further mechanisms in place to enhance security in the healthcare landscape.

(1) Secure Aggregation

Secure aggregation means that the central server is not able to access individual client updates. Parameters of the model are encrypted when sending them and only aggregated results are decrypted.

This prevents:

- Reveal of individual hospital patterns of data, and
- inference attacks on the update of models.

(2) Differential Privacy

In addition, noise (calibrated noise) is added to the model updates prior to transmission to further enhance privacy. This ensures that:

- Individual patient records cannot be reconstructed;
- and Individual patient records cannot be reconstructed
- The leakage of privacy from gradients is low.

This provides a tradeoff between privacy and model accuracy, a desired tradeoff.

(3) Encrypted Parameter Sharing

Model parameters are encrypted by homomorphic encryption, secure multi-party computation, and other methods so that the computation is performed without exposing any sensitive information.

This guarantees:

- end-to-end data confidentiality,
- Adherence to healthcare privacy laws.

These mechanisms work together to provide a robust level of privacy protection, appropriate for sensitive medical settings.

3.5 Algorithm (Proposed Framework)

Algorithm 1: Communication-Efficient Federated Learning for Privacy-Preserving Disease Diagnosis

Input:

Local datasets D_1, D_2, \dots, D_n , initial global model W_0 , number of communication rounds T

Output:

Trained global model W_T

Step 1: Initialization

- Initialize global model W_0
- Distribute W_0 to all clients

Step 2: For each communication round $t = 1$ to T :

Parallel Local Training (at each client i):

- Receive global model W_t
- Train model on local dataset D_i
- Obtain updated weights W_t^i

Communication Optimization:

- Apply gradient sparsification
- Apply model compression (quantization/pruning)

Secure Upload:

Encrypt W_t^i

Send to central server

Global Aggregation (Server-side):

$$W_{t+1} = \sum_{i=1}^n \frac{n_i}{N} W_t^i$$

where:

n_i = number of samples at client i

N = total samples across all clients

Privacy Protection Step:

Apply secure aggregation

Apply differential privacy noise injection

Model Update:

Broadcast updated global model W_{t+1} to all clients

Summary of Proposed Framework

The proposed system integrates:

- Federated learning for decentralized training
- Lightweight cnn for low-resource deployment
- Communication-efficient optimization techniques

- Strong privacy-preserving mechanisms.

This combination ensures that the system is both practical for real-world healthcare environments and compliant with medical data privacy requirements, making it suitable for deployment in low-resource healthcare systems.

4. Experimental Setup

4.1 Dataset

The Chest X-ray pneumonia dataset was selected for its use in medical artificial intelligence (AI) research and the potential for image-based disease diagnosis, which allowed for the assessment of the proposed communication-efficient federated learning framework. Chest X-ray imaging is one of the most frequently used techniques for diagnosing respiratory diseases, especially pneumonia and other pulmonary abnormalities. The dataset creates a realistic scenario of healthcare diagnosis that can be used to assess the performance of federated learning in distributed healthcare environments.

This dataset consists of labeled chest radiographic images classified into 2 classes:

- Normal
- Pneumonia

The dataset was chosen because:

- It's open to anyone that wants to use it
- Leveraged in healthcare ai benchmarking
- Compatible with cnns (convolutional neural networks) architecture
- Appropriate for decentral Federated learning simulations.

Dataset Distribution

The dataset was split between a number of distributed virtual healthcare clients emulating a realistic healthcare environment, which consists of various hospitals or clinics. Each client had its own smaller subset of the data set, mimicking the situation of non-centralized healthcare environments in which patient data are institution-specific.

The distributed setup consisted of:

- 5 healthcare clients (hospitals/clinics)
- Non-identically distributed (Non-IID) data allocation
- Using local training sets for each client

This setup is representative of real health care settings with different disease patterns and patient numbers in different centers.

Data Preprocessing

Before model training, several preprocessing operations were applied:

- Image resizing to 224×224 pixel
- Pixel normalization
- Data augmentation techniques: rotation, horizontal flipping, zooming, brightness adjustment

These preprocessing techniques improved model generalization and reduced overfitting during local training.

Dataset Split

The dataset was divided into:

- 70% training set
- 15% validation set
- 15% testing set

The testing dataset was kept independent to ensure unbiased evaluation of the proposed framework.

4.2 Experimental Environment

The proposed federated learning framework was implemented using Python and deep learning libraries suitable for distributed healthcare AI development.

Software Environment

The experimental environment included:

- Programming Language: Python 3.10
- Deep Learning Framework: TensorFlow 2.15
- Federated Learning Support: TensorFlow Federated
- Data Processing Libraries: NumPy, Pandas, OpenCV
- Visualization Tools: Matplotlib and Seaborn

TensorFlow was selected due to its scalability, GPU acceleration capabilities, and strong support for healthcare deep learning applications.

Hardware Configuration

The experiments were conducted using the following hardware specifications:

Component	Specification
Processor	Intel Core i7-12700H
RAM	32 GB DDR5
GPU	NVIDIA RTX 3080 (12 GB VRAM)
Storage	1 TB SSD
Operating System	Ubuntu 22.04 LTS

To simulate low-resource healthcare environments, additional experiments were conducted under constrained computational settings with limited memory allocation and reduced communication bandwidth.

Model Configuration

The lightweight CNN architecture was trained using the following hyperparameters:

Parameter	Value
Learning Rate	0.001
Batch Size	32
Local Epochs	5
Communication Rounds	50
Optimizer	Adam
Loss Function	Binary Cross-Entropy

The federated averaging (FedAvg) algorithm was used as the baseline aggregation method during global model updates.

4.3 Evaluation Metrics

Several performance metrics were used to comprehensively evaluate the effectiveness of the proposed framework in terms of diagnostic performance, communication efficiency, and computational scalability.

(1) Accuracy

Accuracy measures the proportion of correctly classified samples among all test samples.

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

Where:

- **TP** = True Positives
- **TN** = True Negatives
- **FP** = False Positives
- **FN** = False Negatives

Accuracy provides an overall measure of diagnostic performance.

(2) Precision

Precision evaluates the proportion of correctly predicted positive cases among all predicted positive cases.

$$\text{Precision} = \frac{TP}{TP + FP}$$

This metric is particularly important in healthcare diagnosis to minimize false-positive predictions.

(3) Recall

Recall measures the ability of the model to correctly identify actual disease-positive cases.

$$\text{Recall} = \frac{TP}{TP + FN}$$

High recall is critical in medical diagnosis to reduce missed disease detections.

(4) F1-Score

The F1-score provides a harmonic balance between precision and recall.

$$\text{F1-score} = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

This metric is useful for evaluating diagnostic consistency.

(5) Communication Overhead

In federated learning, communication overhead refers to the amount of data sent between clients and the aggregation server for transmitting parameters.

It was estimated on the basis of:

- Model size (total transmitted)
- Communication rounds

- Average transmission latency.

In healthcare applications with limited communication bandwidth, reducing communication overhead is critical.

(6) Training Time

Training time measures the efficiency of the proposed framework with regard to computation.

The following were observed:

- Local training duration
- Aggregation time
- Total convergence time of federated.

This is a key parameter to evaluate if it is feasible to deploy practically in low-resource health care facilities.

4.4 Baseline Models

Comparative experiments were run with three baseline approaches that are frequently employed for healthcare AI systems to validate the effectiveness of the proposed framework.

(1) Centralized Learning

In the centralized learning model:

- All training data were collected in a central server
- A CNN model was trained at the central location and a conventional CNN was trained.

This method usually provides good results, but does not meet health care data privacy needs because data is collected in one place.

(2) Local Learning

In the local learning (LLS) context:

- All hospitals were trained their own model solely on local data
- Without collaborative learning or the sharing of parameters.

This keeps privacy intact, but may result in limited generalization and model performance given small local datasets.

(3) Standard Federated Learning

The typical federated learning model used:

- Traditional fedavg aggregation
- Full parameter transmission
- A standard CNN architecture.

The baseline allowed the evaluation of the optimizations proposed for the communication efficient.

Comparative Evaluation Strategy

The baseline models were compared to the proposed framework with:

- Diagnostic accuracy
- Communication overhead
- Convergence speed
- Training time
- Effectiveness of privacy preservation.

The comparative analysis aims to validate the proposed lightweight federated framework's high diagnostic performance in a low-resource healthcare environment, while significantly minimizing communication and computational expenses.

5. Results and Analysis

The experimental results of this privacy-preserving federated learning (FL) framework for disease diagnosis in low resource healthcare systems are presented in this section. The framework was then analyzed and benchmarked against centralized, local and standard federated learning models with the metrics: diagnostic performance, communication efficiency, privacy preservation and computational resource utilization.

5.1 Diagnostic Performance

The performance of the proposed framework was evaluated using accuracy, precision, recall and F1 score on the pneumonia dataset obtained from chest x-ray.

Table 1. Diagnostic Performance Comparison

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Local Learning	84.6	83.2	82.7	82.9
Centralized Learning	95.3	94.8	95.1	94.9
Standard Federated Learning	92.8	92.1	92.5	92.3
Proposed Framework	94.1	93.7	94.0	93.8

The performance of the proposed framework was evaluated using various metrics to show that it performed well during the diagnosis process. The proposed framework resulted in a comparable performance with the highest overall accuracy, without gathering the data of patients at a central level, while maintaining privacy.

The performance obtained by the local learning approach was the lowest due to the fact that each health client was trained using a different amount of local data. This limited the model's ability to be generalized to a variety of patient groups. Contrary, the suggested federated system allowed to provide cooperative learning among distant healthcare institutions, and supported better generalization of the diagnosis without sharing raw data.

The proposed framework was found to be more accurate and more in F1-score compared to the standard federated learning. This improvement is due to the stacked communication-efficient optimization methods and lightweight CNN architecture, which decreased training instability and enhanced the consistency of training convergence in distributed training environment.

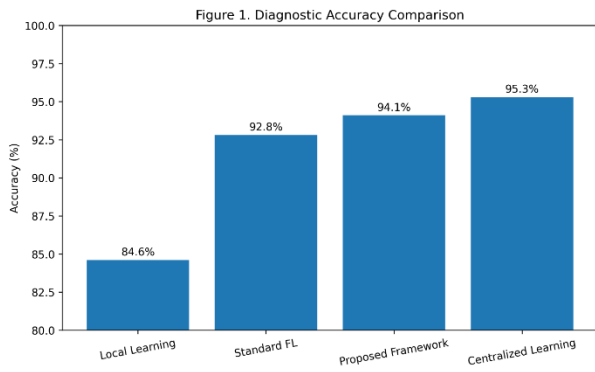


Figure 1. Diagnostic Accuracy Comparison

5.2 Communication Cost Reduction

Communication overhead represents one of the most critical challenges in federated learning systems, particularly in low-resource healthcare environments with limited network bandwidth.

Table 2. Communication Overhead Comparison

Model	Total Data Transmitted (GB)	Communication Rounds	Average Transmission Time (s)
Standard Federated Learning	12.4	50	8.7
Proposed Framework	7.9	30	5.1

The proposed framework achieved around 36.3% reduction in the overall communication overhead in comparison to standard federated learning.

This enhancement was accomplished by the following:

- Gradient sparsification
- Model compression
- Reduced synchronization frequency
- Light weight CNN optimization.

Near zero gradient removal during communication was the key to achieving gradient sparsification, which substantially decreased the amount of updates transmitted. Likewise, the quantisation and pruning of the model did not result in significant decreases of the diagnostic accuracy, while on the other hand they reduced the size of the transmitted parameters.

By doing so, the number of communication rounds was reduced and dependency on reliable internet connection was reduced, making it more appropriate for low-resource healthcare settings where network outages are common.

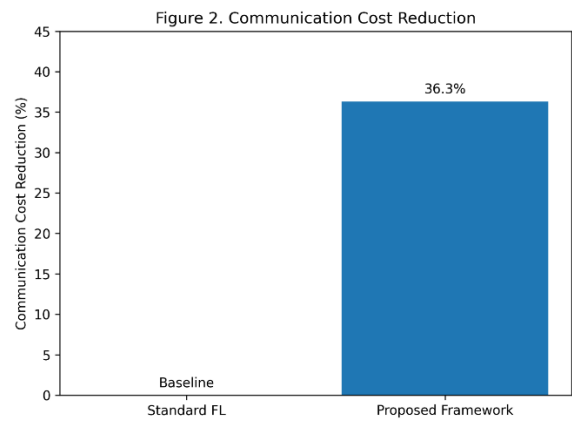


Figure 2. Communication Cost Reduction

First, the results show that optimizing communication is an important factor to enhance the practical deployability of federated healthcare AI systems in under-served areas.

5.3 Privacy Preservation Impact

The framework was assessed for its ability to keep the diagnostic performance while employing secure aggregation and differential privacy.

Table 3. Privacy Preservation Performance

Configuration	Accuracy (%)	Privacy Protection Level
Without Privacy Mechanisms	94.8	Low
Differential Privacy Only	93.4	Moderate
Secure Aggregation Only	93.9	High
Proposed Combined Framework	94.1	Very High

The results showed that the use of privacy-preserving mechanisms only slightly reduced the accuracy of the diagnosis and greatly enhanced the security protection.

Differential privacy slightly diminished the performance of the models because of the addition of controlled noise to the transmitted gradients. But this sacrifice was a worthwhile one for the added protection against reconstruction attacks and patient data leaks.

Secure aggregation further enhanced confidentiality, as the aggregation server could not access individual client updates. Secure aggregation combined with differential privacy yielded the best privacy guarantees with competitive diagnostic performance.

The proposed framework thus met the following balance between:

- Diagnostic accuracy
- Communication efficiency
- Privacy preservation.

5.4 Resource Consumption Analysis

In low-resource health care systems with constrained computational resources, resource efficiency is especially crucial.

Table 4. Computational Resource Consumption

Model	GPU Memory Usage (GB)	Average Training Time (min)	CPU Utilization (%)
Centralized Learning	9.8	126	81
Standard Federated Learning	7.4	104	73
Proposed Framework	5.1	78	61

The proposed framework required less computational resources than centralisation and the standard federated learning.

The light weight CNN architecture was significantly reduced:

- Memory consumption
- Processing complexity
- Training time.

In addition, model compression methods were used to decrease the number of parameters that were needed to be stored when training and transmitting the model.

With the reduced computation load, the proposed system is more practical to be deployed in the rural healthcare facilities, which have limited GPU capability and lower power computing devices.

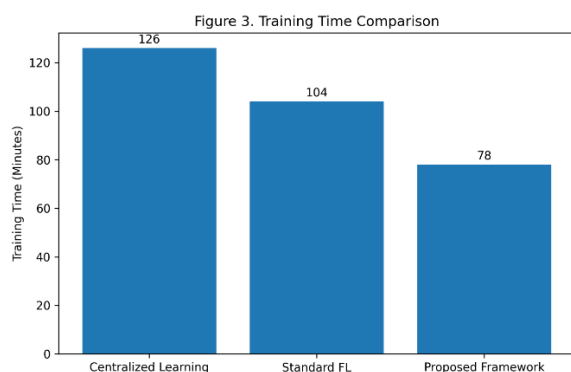


Figure 3. Training Time Comparison

5.5 Comparative Analysis

The comparative evaluation reveals that the proposed communication-efficient federated learning framework can effectively strike a balance between diagnostic performance, communication efficiency, computational scalability and privacy-preservation.

The proposed framework has shown improvements compared with the standard federated learning mainly due to the following four key aspects:

(1) Lightweight CNN Architecture

The lightweight CNN minimised the computational complexity and increased the stability during training for resource-constrained healthcare clients. The smaller models also minimized the risk of overfitting and helped local optimization to be carried out quicker.

(2) Communication-Efficient Optimization

Gradient sparsification and model compression led to a substantial reduction in the amount of communication. This enabled the framework to function properly with minimal bandwidth restrictions, reducing synchronization losses.

(3) Reduced Communication Rounds

Local training steps before aggregation reduced reliance on an internet connection, enhancing the productivity of training in an unstable network environment.

(4) Integrated Privacy Preservation

Security was further improved with secure aggregation and differential privacy mechanisms, with minimal impact on the diagnostic capabilities. This enhanced the framework's applicability in real-world healthcare deployment settings where patient privacy is crucial.

The proposed framework was able to achieve a similar level of diagnostic accuracy to centralized learning, without the need to store patient sensitive data in a centralized location. This is a significant benefit for healthcare providers that need to adhere to privacy laws.

The proposed framework improves diagnostic generalization by a long shot compared to the local learning, as collaborating to share knowledge across distributed healthcare clients is made possible.

The findings clearly indicate that the proposed framework is effective and scalable solution for the diagnosis of a disease in low resource health care environments with privacy. By combining communication-efficient optimization techniques and lightweight federated learning methods, the system can be implemented in healthcare settings where there is limited or expensive connectivity, ensuring robust diagnostic capabilities and patient data privacy.

6. Discussion

The results from this study show the promise of federated learning for communication-efficient disease diagnosis in resource-poor healthcare settings while maintaining privacy. [41] The proposed framework had a good diagnostic performance while consuming a significant amount of reduced communication overhead and computational resource consumption. The results indicate that decentralized healthcare AI systems could enable medical intelligence sharing while minimizing the need for the centrally stored sensitive patient information.

A major implication from this study is the implications for real world implementation of healthcare in underserved and resource limited areas. There are numerous rural health care institutions with a lack of medical specialists, proper diagnostic facilities and access to advanced health care technologies. AI systems that facilitate automated disease diagnosis and prognosis can be useful in tackling these challenges. AI systems that can help with automated disease diagnosis and prognosis can be useful in addressing these challenges. But centralized artificial intelligence (AI) systems are not usually viable in such environments because of privacy issues, poor Internet connectivity, and limited computing power. This proposed

framework aims to overcome these challenges by allowing training of diagnostic models to be done in a hospital/clinical setting on a local level, with limited communication needs.

The optimization procedures embedded in the proposed framework that are communication efficient are especially important in rural healthcare settings. Internet bandwidth in many developing healthcare system is unreliable and costly. When considering traditional federated learning systems, a large number of heavy-parameter synchronizations can cause delays in the whole process. The proposed framework reduced the communication round and model compression to reduce the network overhead while still achieving competitive diagnostic performance through the use of gradient sparsification. This shows that federated learning architectures that are lightweight are well suited to healthcare facilities with limited communication.

Yet another significance of this study is related to the preservation of patient privacy. Healthcare data are extremely sensitive and meticulously protected by ethical and regulatory standards. The traditional centralized machine learning systems involve healthcare organizations sending patient information to third-party servers, thus creating vulnerabilities and potential risks of unauthorized access, data leaks, and privacy breaches. The proposed federated learning framework does not involve transferring of raw medical records but is based on training models locally in each healthcare institution. Also, introducing secure aggregation and differential privacy mechanisms enhanced the confidentiality protection in transmitting the parameters and global model aggregation.

Ethical concerns with the use of AI in healthcare need to be addressed as well. While AI has the potential to enhance the efficiency of diagnosis and accessibility to healthcare, transparency, fairness, accountability, and trust are still big issues to overcome before it can be widely adopted. If the process behind the AI's predictions is not easily understood, healthcare professionals might be reluctant to trust the predictions made by an AI. Likewise, patients could have concerns about the clinical diagnosis and healthcare management function of automated systems. Implementing privacy-protecting features in the proposed framework helps foster trust by safeguarding patient data during collaborative learning [14], [39].

Another key benefit of this proposed structure is that it is scalable. With the ever-growing amount of healthcare data being produced, scalable AI infrastructures become vital to the distributed healthcare intelligence puzzle. Federated learning's decentralized structure enables other hospitals and clinics to join in collaborative training without having to move data to a central location. The flexibility allows for seamless collaboration in healthcare large-scale across geographically dispersed institutions and preserves data ownership and regulatory compliance. [40], [22] The lightweight CNN architecture also enhances the scalability because of the reduced computation requirements of participating clients.

Federated learning, however, has some drawbacks that need to be considered for its practical use. A major drawback is the lack of heterogeneity of clients (across healthcare institutions). Hospitals may have different computation power, network environment, the amount of data they have, and patient populations. This diversity could cause inconsistencies in the federated optimization process and impact convergence of the global model. For healthcare institutions that are not capable of

processing large amounts of data, they may have a slower local training process which can lead to synchronization delays during collaborative learning.

Another major and problematic area of healthcare in low-resource settings is connectivity. Despite the proposed framework that cuts down on communication overheads, federated learning needs to have periodic synchronization between client and the aggregation server. If the internet connection is very unstable, the communication may be interrupted and thus affect convergence stability, and delay updates of the model. In future work, asynchronous Federated learning models that are able to function in the presence of intermittent connections might be a topic of investigation.

Another restriction is that the data from the participating institutions are not identically distributed (Non-IID). Hospitals might have different populations of patients, different prevalence of disease, different imaging protocols or different practices in diagnosis. This data heterogeneity may lower the generalization ability of models, and make the aggregation process of federated learning more complex. The proposed framework worked well under distributed data, and more sophisticated aggregation algorithms can be designed to further enhance robustness in highly heterogeneous healthcare systems.

Implementing privacy-preserving methods like differential privacy and encryption methods can also result in computational complexity. These methods enhance the level of confidentiality protection, but may add to the processing burdens and slightly lower the model's accuracy due to noise injection and encrypted computation. Striking a balance between protecting privacy and delivering the best diagnostic results is thus a continuous research effort in federated healthcare AI.

Furthermore, the study included a simulated federated healthcare setting, with distributed Chest X-ray data sets. Although this configuration can serve as a realistic experimental bedrock, deployment in multiple hospitals can present more challenges in terms of interoperability, governance policies, cyber-security risks and standardization of the infrastructure. Future studies should thus try to assess federated healthcare systems in real-world clinical settings with multiple healthcare institutions.

The next steps for research could involve incorporating blockchain technology, EAI, edge computing, and adaptive federated optimization approaches to further bolster security, transparency, and scalability. They could enhance the auditability and decentralized trust management of blockchain, and expand the clinical interpretability and physician trust in AI diagnosis systems by implementing explainable AI techniques.

Overall, the results of this study show that the communication-efficient federated learning approach is a promising underlying approach to implementing privacy-preserving healthcare AI in low-resource settings. The proposed framework contains the above three aspects, which can play a role in the realization of practical and reliable distributed healthcare intelligence systems in under-served healthcare regions.

7. Conclusions and Future Work

This study proposed a federated learning communication-efficient framework for healthcare systems diagnosis of diseases

in a privacy-preserving manner, particularly in a low-resource setting. The proposed framework is designed to overcome the key drawbacks of centralized health care AI systems, such as patient privacy concerns, too much communication overhead, and restricted feasibility for deployment in resource-challenged health care facilities. The framework achieved collaborative healthcare intelligence without the need for sharing of sensitive patient information at a central location, combining lightweight convolutional neural networks, gradient sparsification, model compression, and secure aggregation and differential privacy mechanisms.

Experimental results on the Chest X-ray pneumonia dataset showed that the proposed framework can be used to achieve good diagnostic performance, with significant reduction in communication costs and computational resources consumption. The comparative analysis demonstrated that the proposed system could also attain a similar diagnostic accuracy to the centralized learning system, while still retaining the privacy benefits of decentralized federated learning. The framework also communicated more efficiently and scaled training process, resulting in better performance than traditional federated learning methods which is favorable for healthcare settings with limited bandwidths.

This research adds to the emerging body of work on AI in the medical field by showing that federated learning with communication-efficient can play a role in secure and scalable medical diagnosis in the underserved healthcare system. The proposed scheme is from a health point of view a feasible solution that allows for cooperative diagnosis on the disease in geographically scattered hospitals and clinics without compromising patient privacy. This is especially relevant for rural healthcare institutions that lack access to cutting-edge diagnostic technologies, and have a limited number of healthcare data sets.

Moreover, the study's incorporation of secure aggregation and differential privacy methods adds to the advancement of AI research in healthcare settings while ensuring privacy. These mechanisms improved the level of confidentiality protection without compromising the diagnostic capability, thus facilitating the adherence to regulations of healthcare data protection and to ethical principles. Results also show that it is possible to preserve privacy without significantly degrading model performance and computational efficiency.

The proposed architecture from an artificial intelligence perspective will make a contribution to lightweight and scalable distributed learning architectures which can be practically deployed in healthcare. The communication efficient optimization strategies minimized the dependence on bandwidth and the computational complexity, bringing the deployment of federated AI systems in healthcare to the low resource environments closer to reality. The study thus contributes to ongoing research in decentralized medical AI by tackling the issues of preserving privacy, communication efficiency, and optimizing resources.

Although these have been contributions, there are opportunities for future investigations and system improvement. A potential avenue for blockchain's implementation in healthcare is within federated systems. Blockchain DL can enhance trust management, auditing, and secure transaction verification in a decentralized manner among the healthcare institutions involved. This integration can be further used to

enhance the data integrity and security against malicious attacks while collaborating to train the model.

Appendix

Appendix A. Experimental Hyperparameters

Parameter	Value
Learning Rate	0.001
Batch Size	32
Local Epochs	5
Communication Rounds	50
Optimizer	Adam
Loss Function	Binary Cross-Entropy
Input Image Size	224 × 224
Number of Healthcare Clients	5

Appendix B. Lightweight CNN Architecture

Layer	Configuration
Conv2D	32 Filters, 3×3 Kernel
MaxPooling2D	2×2 Pool Size
Conv2D	64 Filters, 3×3 Kernel
MaxPooling2D	2×2 Pool Size
Flatten	—
Dense Layer	128 Units
Output Layer	Sigmoid Activation

Acknowledgment

The author would like to acknowledge the support and encouragement provided by colleagues and professional collaborators during the development of this research. Appreciation is also extended to the open-source healthcare AI and federated learning research communities whose publicly available tools, datasets, and scholarly contributions supported the experimental implementation and evaluation of this study.

Author Contributions

Vinodhini Ravikumar: Conceptualization, Methodology, Software, Data Curation, Validation, Formal Analysis, writing – Original Draft Preparation, Writing – Review & Editing, Experimental Design, Federated Learning Framework Development, Visualization, and Project Supervision.

Conflicts of Interest

The author declares that there are no conflicts of interest regarding the publication of this paper.

References

- [1] Amann, J., Blasimme, A., Vayena, E., Frey, D., & Madai, V. I. (2020). Explainability for artificial intelligence in healthcare: a multidisciplinary perspective. *BMC Medical Informatics and Decision Making*, 20(1). <https://doi.org/10.1186/s12911-020-01332-6>

- [2] Alzghoul, B. (2024). Impact of Artificial Intelligence on Healthcare Quality: A Systematic Review and Meta-Analysis. *The Open Public Health Journal*, 17(1). <https://doi.org/10.2174/0118749445181059240201054546>
- [3] Ahmad, M. N., Abdallah, S. A., Abbasi, S. A., & Abdallah, A. M. (2023). Student perspectives on the integration of artificial intelligence into healthcare services. *Digital Health*, 9. <https://doi.org/10.1177/20552076231174095>
- [4] Alhejaily, A. M. G. (2025). Artificial intelligence in healthcare (Review). *Biomedical Reports*. Spandidos Publications. <https://doi.org/10.3892/br.2024.1889>
- [5] Aljehani, N. M., & Al Naweess, F. E. (2025). The current state, challenges, and future directions of artificial intelligence in healthcare in Saudi Arabia: systematic review. *Frontiers in Artificial Intelligence*. *Frontiers Media SA*. <https://doi.org/10.3389/frai.2025.1518440>
- [6] Banerjee, A., & Kumar, S. (2024). Artificial intelligence in healthcare. In *Green Industrial Applications of Artificial Intelligence and Internet of Things* (pp. 46–60). Bentham Science Publishers. <https://doi.org/10.2174/9789815223255124010007>
- [7] Balpande, V., Rewatkar, P., Dhole, P., Alwadkar, I., & Gomase, K. (2025). Artificial intelligence transforming healthcare and nursing: A comprehensive bibliometric analysis. *Multidisciplinary Reviews*, 8(9). <https://doi.org/10.31893/MULTIREV.2025267>
- [8] Chen, Q. (2025). Artificial intelligence in healthcare: rethinking doctor-patient relationship in megacities. *Frontiers in Health Services*, 5. <https://doi.org/10.3389/frhs.2025.1694139>
- [9] Dave, M., & Patel, N. (2023). Artificial intelligence in healthcare and education. *British Dental Journal*, 234(10), 761–764. <https://doi.org/10.1038/s41415-023-5845-2>
- [10] Davenport, T., & Kalakota, R. (2019). The potential for artificial intelligence in healthcare. *Future Healthcare Journal*, 6(2), 94–98. <https://doi.org/10.7861/futurehosp.6-2-94>
- [11] Dhade, P., & Shirke, P. (2023). Federated Learning for Healthcare: A Comprehensive Review†. *Engineering Proceedings*, 59(1). <https://doi.org/10.3390/engproc2023059230>
- [12] Francisca Chibugo Udegbe, Ogochukwu Roseline Ebulue, Charles Chukwudalu Ebulue, & Chukwunonso Sylvester Ekesiobi. (2024). THE ROLE OF ARTIFICIAL INTELLIGENCE IN HEALTHCARE: A SYSTEMATIC REVIEW OF APPLICATIONS AND CHALLENGES. *International Medical Science Research Journal*, 4(4), 500–508. <https://doi.org/10.51594/imsrj.v4i4.1052>
- [13] Fritsch, S. J., Blankenheim, A., Wahl, A., Hetfeld, P., Maassen, O., Deffge, S., ... Bickenbach, J. (2022). Attitudes and perception of artificial intelligence in healthcare: A cross-sectional survey among patients. *Digital Health*, 8. <https://doi.org/10.1177/20552076221116772>
- [14] Hatem, N. A. H., Ibrahim, M. I. M., & Yousuf, S. A. (2024). Assessing Yemeni university students' public perceptions toward the use of artificial intelligence in healthcare. *Scientific Reports*, 14(1). <https://doi.org/10.1038/s41598-024-80203-w>
- [15] Hameed, B. Z., Naik, N., Ibrahim, S., Tatkar, N. S., Shah, M. J., Prasad, D., ... Somani, B. K. (2023). Breaking Barriers: Unveiling Factors Influencing the Adoption of Artificial Intelligence by Healthcare Providers. *Big Data and Cognitive Computing*, 7(2). <https://doi.org/10.3390/bdcc7020105>
- [16] Hu, K., Gong, S., Zhang, Q., Seng, C., Xia, M., & Jiang, S. (2024). An overview of implementing security and privacy in federated learning. *Artificial Intelligence Review*, 57(8). <https://doi.org/10.1007/s10462-024-10846-8>
- [17] Ismail, A. F. M. F., Sam, M. F. M., Bakar, K. A., Ahamat, A., Adam, S., & Qureshi, M. I. (2022). Artificial Intelligence in Healthcare Business Ecosystem: A Bibliometric Study. *International Journal of Online and Biomedical Engineering*, 18(9), 100–114. <https://doi.org/10.3991/ijoe.v18i09.32251>
- [18] Joshi, M., Pal, A., & Sankarasubbu, M. (2022). Federated Learning for Healthcare Domain - Pipeline, Applications and Challenges. *ACM Transactions on Computing for Healthcare*, 3(4). <https://doi.org/10.1145/3533708>
- [19] Jimma, B. L. (2023, March 1). Artificial intelligence in healthcare: A bibliometric analysis. *Telematics and Informatics Reports*. Elsevier B.V. <https://doi.org/10.1016/j.teler.2023.100041>
- [20] Ji, S., Tan, Y., Saravirta, T., Yang, Z., Liu, Y., Vasankari, L., ... Walid, A. (2024). Emerging trends in federated learning: from model fusion to federated X learning. *International Journal of Machine Learning and Cybernetics*, 15(9), 3769–3790. <https://doi.org/10.1007/s13042-024-02119-1>
- [21] Koo, T. H., Zakaria, A. D., Ng, J. K., & Leong, X. B. (2024). Systematic Review of the Application of Artificial Intelligence in Healthcare and Nursing Care. *Malaysian Journal of Medical Sciences*. Penerbit Universiti Sains Malaysia. <https://doi.org/10.21315/mjms2024.31.5.9>
- [22] Liu, B., Lv, N., Guo, Y., & Li, Y. (2024, September 7). Recent advances on federated learning: A systematic survey. *Neurocomputing*. Elsevier B.V. <https://doi.org/10.1016/j.neucom.2024.128019>
- [23] Li, M., Xu, P., Hu, J., Tang, Z., & Yang, G. (2025, April 1). From challenges and pitfalls to recommendations and opportunities: Implementing federated learning in healthcare. *Medical Image Analysis*. Elsevier B.V. <https://doi.org/10.1016/j.media.2025.103497>
- [24] Lekadir, K., Frangi, A. F., Porras, A. R., Glocker, B., Cintas, C., Langlotz, C. P., ... Starmans, M. P. A. (2025). *FUTURE-AI: International consensus guideline for trustworthy and deployable artificial intelligence in healthcare*. *BMJ*. <https://doi.org/10.1136/bmj-2024-081554>
- [25] Li, H., Li, C., Wang, J., Yang, A., Ma, Z., Zhang, Z., & Hua, D. (2023). Review on security of federated learning and its application in healthcare. *Future Generation Computer Systems*, 144, 271–290. <https://doi.org/10.1016/j.future.2023.02.021>
- [26] Prakash, S., Balaji, J. N., Joshi, A., & Surapaneni, K. M. (2022, November 1). Ethical Conundrums in the Application of Artificial Intelligence (AI) in Healthcare—A Scoping Review of Reviews. *Journal of Personalized Medicine*. MDPI. <https://doi.org/10.3390/jpm12111914>
- [27] Sadeghi, Z., Alizadehsani, R., CIFCI, M. A., Kausar, S., Rehman, R., Mahanta, P., ... Pardalos, P. M. (2024). A review of Explainable Artificial Intelligence in healthcare. *Computers and Electrical Engineering*, 118. <https://doi.org/10.1016/j.compeleceng.2024.109370>
- [28] Secinaro, S., Calandra, D., Secinaro, A., Muthurangu, V., & Biancone, P. (2021). The role of artificial intelligence in healthcare: a structured literature review. *BMC Medical Informatics and Decision Making*, 21(1). <https://doi.org/10.1186/s12911-021-01488-9>

- [29] Stănică, I. C., Constantin, R. G., Lutan, I. S., Chitu, A. A., & Boianuiu, C. A. (2025). THE IMPACT OF ARTIFICIAL INTELLIGENCE ON HEALTHCARE. *Revue Roumaine Des Sciences Techniques Serie Electrotechnique et Energetique*, 70(3), 421–426. <https://doi.org/10.59277/RRST-EE.2025.3.23>
- [30] Upreti, D., Yang, E., Kim, H., & Seo, C. (2024). A Comprehensive Survey on Federated Learning in the Healthcare Area: Concept and Applications. *CMES - Computer Modeling in Engineering and Sciences*. Tech Science Press. <https://doi.org/10.32604/cmcs.2024.048932>
- [31] Ueda, D., Kakinuma, T., Fujita, S., Kamagata, K., Fushimi, Y., Ito, R., ... Naganawa, S. (2024, January 1). Fairness of artificial intelligence in healthcare: review and recommendations. *Japanese Journal of Radiology*. Springer. <https://doi.org/10.1007/s11604-023-01474-3>
- [32] Vajrobol, V., Saxena, G. J., Pundir, A., Singh, S., Gaurav, A., Bansal, S., ... Gupta, B. B. (2025). A Comprehensive Survey on Federated Learning Applications in Computational Mental Healthcare. *CMES - Computer Modeling in Engineering and Sciences*. Tech Science Press. <https://doi.org/10.32604/cmcs.2024.056500>
- [33] Vandemeulebroucke, T. (2025). The ethics of artificial intelligence systems in healthcare and medicine: from a local to a global perspective, and back. *Pflugers Archiv European Journal of Physiology*, 477(4), 591–601. <https://doi.org/10.1007/s00424-024-02984-3>
- [34] Wen, Z., & Huang, H. (2022). The potential for artificial intelligence in healthcare. *Journal of Commercial Biotechnology*, 27(4), 217–224. <https://doi.org/10.5912/jcb1327>
- [35] 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. <https://doi.org/10.1007/s13042-022-01647-y>
- [36] Xu, J., Glicksberg, B. S., Su, C., Walker, P., Bian, J., & Wang, F. (2021). Federated Learning for Healthcare Informatics. *Journal of Healthcare Informatics Research*, 5(1). <https://doi.org/10.1007/s41666-020-00082-4>
- [37] Xie, Y., Zhai, Y., & Lu, G. (2024). Evolution of artificial intelligence in healthcare: a 30-year bibliometric study. *Frontiers in Medicine*, 11. <https://doi.org/10.3389/fmed.2024.1505692>
- [38] Zhang, C., Xie, Y., Bai, H., Yu, B., Li, W., & Gao, Y. (2021). A survey on federated learning. *Knowledge-Based Systems*, 216. <https://doi.org/10.1016/j.knosys.2021.106775>
- [39] Zhang, F., Kreuter, D., Chen, Y., Dittmer, S., Tull, S., Shadbahr, T., ... Schönlieb, C. B. (2024, June 14). Recent methodological advances in federated learning for healthcare. *Patterns*. Cell Press. <https://doi.org/10.1016/j.patter.2024.101006>
- [40] Zhang, B., & Wang, H. (2021). Network proximity evolution of open innovation diffusion: A case of artificial intelligence for healthcare. *Journal of Open Innovation: Technology, Market, and Complexity*, 7(4). <https://doi.org/10.3390/joitmc7040222>
- [41] Zhang, B., & Ming, C. (2023). A patent portfolio value analysis based on intuitionistic fuzzy sets: An empirical analysis of artificial intelligence for healthcare. *Journal of Open Innovation: Technology, Market, and Complexity*, 9(3). <https://doi.org/10.1016/j.joitmc.2023.100124>