Journal of Emerging Trends in Computer Science and Applications (JETCSA), Volume
1 (Issue 1: January-April, 2025): Pages 1-17; Published on: May 05, 2025.

ISSN: XXXX-XXXX Awaiting for Approval (Online)
Journal of Emerging Trends in Computer Science and Applications(JETCSA)
Contents available at: *https://www.swamivivekanandauniversity.ac.in/jetcse/*

Language Models for Federated Continuous Learning for Privacy-Preserving Medical Diagnosis

Mrs. Sangita Bose[1,*], Dr. Romit S. Beed[2]

[1]Department of Computer Science,Swami Vivekananda University,Barrackpore,WB,India; *sangitab@svu.ac.in*
[2]Department of Computer Science, St. Xavier's College (Autonomous), Kolkata, India

Abstract

Artificial intelligence is quickly being used in healthcare to help with Artificial intelligence is largely being exploited in healthcare for supporting diagnosis through the analysis of images and reports. It requires an enormous amount of data to train trustworthy diagnosis models. However, sharing sensitive patient information between hospitals or even separate research institutes raises major privacy concerns. Federated learning saves the day by allowing various universities to collaborate to train models without sharing raw data. In real-world clinical contexts, models require not simply pooled learning, but also the ability to learn about new diseases, imaging modalities, and how medical knowledge evolves over time. This necessitates federated constant learning, which enables models to maintain existing knowledge while progressively gaining new skills. We propose employing vision-language models (VLMs) in a federated continual learning framework to diagnose medical conditions. VLMs will combine visual medical data (X-rays and MRIs) with written data (clinical notes and diagnostic labels). By merging these two types of information, our technique improves understanding and performance across a wide range of medical tasks. We integrate strategies for continual learning to ensure that important medical information from the past is retained while adjusting to current circumstances. This method allows us to retain what we have already learned. Our strategy promotes privacy-preserving distributed training among medical institutions while tackling novel diagnostic issues. It enhances diagnostic accuracy and consistency while also fostering trust and collaboration in delicate healthcare situations. This study demonstrates the possibility for scalable, secure diagnostic systems that adapt to the changing nature of healthcare using federated learning, continuous learning, and vision-language modeling.

Keywords: *Federated Learning, Continual Learning, Vision–Language Models, Medical Diagnosis, Privacy Preservation*

1. Introduction

Artificial intelligence (AI) has recently transformed the way clinicians make medical diagnosis. Machine learning (ML) and vision-language models (VLMs) are two new methodologies that integrate images and text. These new models generate technology that can rethink procedures and perform scans from several sources, such as medical notes or X-rays, rather than depending simply on medical data. This feature provides much improved assistance for any diagnostic process. For example, VLMs can

*Author for correspondence

generate X-ray reports that describe the symptoms associated with the images and assist clinicians in distinguishing between similar health conditions.

While the potential of these models in health care is intriguing, there are three key difficulties to their implementation: preserving patient privacy, managing various types of data, and assisting the models in learning quickly. Typically, training AI models necessitates concurrent access to data from multiple hospitals. This is especially challenging in Europe, where strong legislation such as the GDPR (General Data Protection Regulation) apply. Similarly, in the United States, transferring personal health information between hospitals is subject to HIPAA restrictions. Even when sharing is permissible, other considerations, such as obtaining patient consent and avoiding data leaks, complicate matters. This demonstrates the necessity for collaborative solutions that use private data.Federated Learning (FL) could help. In this way, different places, like hospitals, can train a model together without sharing any raw patient data. Instead, only model parameters or gradients are shared with a central server, while sensitive data remains securely stored at each site [5]. This not only reduces privacy risks and ensures compliance with legal requirements but also leverages the collective expertise of diverse medical centers spread across regions. Early use of FL in medical imaging has shown its viability for applications like tumor segmentation, disease classification, and digital pathology examination [6,7]. However, federated learning alone is not enough to address the changing needs of clinical practice. Medical knowledge is not fixed: diagnostic recommendations change, new illnesses arise, and imaging modalities improve with time. Accordingly, AI systems in the field of healthcare should be able to learn continuously, i.e., update themselves with fresh tasks and data distributions without forgetting the knowledge gained previously.

This condition brings forth the framework of Federated Continual Learning (FCL). FCL integrates the ideas of FL with CL, thus allowing distributed AI systems to learn incrementally over time without experiencing catastrophic forgetting [8,9]. In medical diagnosis, catastrophic forgetting is an extremely dangerous condition: a model which forgets to recognize pneumonia when it is trained to recognize COVID-19 would become clinically unsafe. CL methods, including regularization-based methods, memory replay, and parameter isolation, have been suggested to mitigate forgetting in centralized environments [10]. Using these methods in federated environments is not easy because it has to deal with uneven data, limits on communication, and different computing resources at medical centers [11].

With Vision–Language Models, the FCL approach offers new opportunities to facilitate improved medical diagnosis. On occasion, doctors must review images and read text concerning it, such as notes from other doctors or reports about the health of a patient. Jonquet et al [12]) by integrating the information contained in both types of data, VLMs are capable of providing more informative clues for diagnosis than systems that rely on a single data type. For instance, a VLM could learn to interpret mammogram images from one hospital, while also learning from text reports from another at the same time, benefiting both tasks without compromising privacy. Also, as new imaging tools like 3D ultrasounds and new MRIs emerge, the VLM needs regular updates to remain useful in healthcare.

> Using VLMs in FCL systems presents challenges in research and engineering. First, large VLMs like CLIP or BLIP-2 complicate communication when they work together. We need to find ways to reduce data size without losing quality. Second, different data types from various sources create inconsistencies; for example, one hospital may mostly have X-ray images, while another has mainly text files with few images, or even only text files. A reliable method to merge and align this data is essential to ensure our models work consistently across both settings. Third, VLMs should operate correctly together over an extended period. When we update them, we must ensure that the updates keep the images visually appealing while making connections between text and images when necessary. Otherwise, physicians may find it confusing and unhelpful.

- • Ethical and social ramifications should be considered in addition to technical difficulties. Federated continuous VLMs for medical diagnosis present questions about fairness, accountability, and openness. If we do not remove biases in medical datasets, they may be transferred over into federated models, resulting in inequitable diagnoses across different populations [16]. Similarly, continual updates might introduce idea drift, as evidenced by changes in illness incidence rates or updated clinical practice guidelines, which can alter our perception of new facts. As a result, while implementing these technologies in healthcare, it is critical to provide sufficient monitoring, transparency, and medical ethics [17]. To address such issues, this book presents a novel strategy for combining different learning modalities with Vision-Language Models (FCL-VLMs). The salient points are: It brings together federated learning and ongoing learning methods to help protect privacy in medical diagnosis across different healthcare institutions.
- It introduces new ways to avoid losing important past knowledge in Vision–Language Models, ensuring they remain accurate for past tasks while learning new ones.
- It tests the method on well-known clinical data like chest X-rays and medical records to show that it works well and is strong.

By linking these learning methods with Vision–Language Models, this work aims to improve privacy in medical AI. The framework also emphasizes gaining validity in healthcare and adhering to ethical principles so that there a safe and flexible medical AI system.

2. Literature Review

In recent years, individuals have become increasingly aware of how federated learning, ongoing learni ng, and vision-language models operate in medicine.
This section examines the relevant literature in three major areas:
(i) Federated Learning in Medical AI, (ii) Continuous Learning in Medical Diagnosis, and (iii) Vision -Language Models for Medicine.

2.1. Federated Learning in Medical AI

In recent years, more people have become aware of how federated learning, continuous learning, and vision-language models operate in medicine. This section examines notable literature in three major areas: (i) Federated Learning in Medical AI; (ii) Continuous Learning in Medical Diagnosis; and (iii) Vision-Language Models in Medicine. In medical health research, Sheller et al. [2] shown that FL could be used to separate brain tumours across institutions and performed similarly to standard methods. Kaissis et al. [3] and others examined current uses of FL in medical imaging and concluded that it complies with privacy regulations such as HIPAA and GDPR. Dou et al. [4] proposed methods for improving FL by aggregating models from centers in scenarios when the data is not equally distributed across. FL in medical diagnosis presents challenges. Zhao et al. [5] demonstrated that unbalanced client data yield inferior performance of FL. Li et al. [6] suggested various algorithms depending on differing types of networks, corresponding to patterns of improvement performance despite varied client contribution. These findings show that while FL helps keep data private, using it in healthcare will need to overcome issues like varied data, slow internet, and scalability.

2.2. Continual Learning for Medical Diagnosis

Medical diagnosis demands models with the ability to fit dynamic data streams like novel imaging modalities, disease variants, or diagnostic protocols. Repeated retraining of traditional deep learning models on new tasks is affected by catastrophic forgetting—a process where previously acquired knowledge is erased during adaptation [7].

Parisi et al. [8] presented a broad overview of CL methods, sorting them into regularization-based methods, replay memory, and isolation of parameters. Lopez-Paz and Ranzato [9] introduced Gradient Episodic Memory (GEM), which limits the gradient updates to maintain performance on previous tasks. Kirkpatrick et al. [10] presented Elastic Weight Consolidation (EWC), which penalizes the modification of significant parameters, minimizing forgetting.

In medicine, Ozbulak et al. [11] demonstrated the potential of CL for classification in histopathology, where models are required to learn from novel tissue types. Masana et al. [12] pushed the limits of CL by benchmarking in class-incremental scenarios, giving us glimpses into the stability of algorithms over changing datasets.

2.3 Vision–Language Models in Healthcare

Multimodal learning, especially Vision–Language Models (VLMs), has seen growing interest in healthcare of late. *Wang et al.* [13] presented TieNet, a chest X-ray classification and report generation model that learns jointly from images and text radiology reports. This paper illustrated the benefit of fusion of modalities for interpretability and decision support.

Inspired by general-domain models like CLIP [14], researchers have adapted VLMs to the medical domain. Zhang *et al.* [15] proposed a contrastive vision–language pre-training approach for medical applications, showing improved performance on chest X-ray benchmarks. Li *et al.* [16] introduced a medical vision–language model capable of report generation and cross-modal retrieval, emphasizing clinical interpretability.

Despite their promise, VLMs are resource-intensive and typically require centralized training on large-scale multimodal datasets. Johnson *et al.* [17] developed the MIMIC-CXR dataset, a large collection of chest radiographs paired with reports, which has become a benchmark for medical VLMs. However, the centralized use of such dataset's conflicts with privacy concerns, motivating the integration of VLMs with FL and continual learning frameworks.

2.4. Towards Federated Continual Vision–Language Learning

While significant research exists on FL, CL, and VLMs individually, their integration remains underexplored in healthcare. Some preliminary efforts have begun bridging these domains. Wu *et al.* [18] investigated federated multimodal learning, enabling distributed training across text and image data. Similarly, Yang *et al.* [19] introduced FL methods for multimodal EHRs, highlighting cross-site knowledge sharing without data leakage.

In continual learning, Nguyen *et al.* [20] proposed federated continual frameworks to address dynamic task distributions, but their experiments were limited to unimodal datasets. To date, no unified framework robustly addresses federated continual learning with VLMs for medical diagnosis, particularly under non-IID multimodal settings.

2.5. Research Gap

From the literature, three critical gaps emerge:

i. Privacy-preserving VLM training: Existing medical VLMs rely on centralized multimodal datasets, raising privacy concerns. FL offers a solution, but communication efficiency for large models is underexplored.

ii. Continual multimodal adaptation: Most CL methods are validated on unimodal vision tasks. Their extension to VLMs, which require maintaining cross-modal alignment, is still limited.

iii. Federated Continual Integration: While FL and CL have been separately studied, their joint integration for multimodal medical diagnosis is largely absent. Addressing this gap is crucial to ensure adaptive, scalable, and ethical medical AI.

3. Methodology

This study employs a quantitative, experimental design to evaluate a novel framework for Federated Continual Learning (FCL) with Vision-Language Models (VLMs). The primary objective is to demonstrate the feasibility and efficacy of this approach for privacy-preserving medical diagnosis, specifically in scenarios where patient data is distributed across multiple, independent healthcare institutions. The experiment is designed to address key challenges, including data privacy, catastrophic forgetting of previously learned medical conditions, and data heterogeneity across different client sites.

3.1 System Architecture

The proposed system follows a client-server architecture, typical of federated learning. A single **central server** orchestrates the training process, while multiple **client nodes** (representing individual hospitals or clinics) perform local training on their private datasets.

3.1.1 The Central Server

The central server is responsible for model aggregation and distribution. It does not store or have access to any raw patient data. Its functions include:

- **Global Model Initialization:** A pre-trained VLM is initialized and sent to all participating clients at the beginning of the training process.
- **Update Aggregation:** The server receives model updates (weights or gradients) from the clients and aggregates them into a new global model. The aggregation is performed using the **Federated Averaging (FedAvg)** algorithm, weighted by the number of training samples at each client.
- **Global Model Distribution:** The newly aggregated global model is then broadcast back to the clients for the next training round.

3.1.2 The Client-Side

Each client node operates autonomously on its private, non-IID (non-independently and identically distributed) medical data. The core of the client-side methodology is the integration of continual learning with the federated process.

- **Local Data:** Each client dataset consists of pairs of medical images (e.g., chest X-rays) and corresponding clinical text (e.g., radiology reports).
- **Continual Learning Mechanism:** To prevent catastrophic forgetting, each client uses an **Experience Replay (ER)** approach. A small, fixed-size memory buffer stores a subset of data from previous diagnostic tasks (e.g., different diseases). During each local training round, the model is trained on a combination of new data and a small batch of data from the memory buffer.

- **Local Training:** The local VLM is trained for a specified number of epochs using a standard optimization algorithm (e.g., Adam or SGD). The model learns from its new data while being periodically reminded of past knowledge through the experience replay buffer.

3.2 Model Architecture and Implementation

This section details the specific components of the Vision-Language Model and the overall implementation workflow.

3.2.1. Vision-Language Model (VLM)

The VLM architecture is a multi-modal neural network designed to fuse information from both image and text inputs.

- Image Encoder: A pre-trained convolutional neural network (e.g., ResNet-50) is used to extract features from medical images. The final classification layer is removed, and the output of the penultimate layer serves as the image feature vector.
- • Text Encoder: A transformer-based model, like BERT or BioBERT, processes clinical text. The result is a study on context embedding for the full text.
- • The fusion module combines picture and text feature vectors and passes them via fully connected layers. This lesson introduces a combined representation that encapsulates the link between visual discoveries and textual descriptions.
- Classification Head: A final softmax layer outputs the probabilities for a set of predefined medical diagnoses.

3.2.2. Implementation Workflow

The practical implementation involves several rounds of back-and-forth communication between the central server and the clients. The process is as follows:

1. **Central Server Initialization:** The global model weights (W_0) are sent to all N clients.
2. **Client-Side Local Training (Round t):**
    - Each client k∈ {1..., N} receives the global model (W_t).
    - Client k updates its local memory buffer with a small number of new data samples.
    - Client k trains its local model (w_k) for E epochs on a combination of its local dataset and a batch of data from its memory buffer. The loss function for local training is given by:

$$L_k = \sum_{i \in \text{New Data}_k} \mathcal{L}(\text{model}(x_i), y_i) + \lambda \sum_{j \in \text{Memory}_k} \mathcal{L}(\text{model}(x_j), y_j)$$

where L is the cross-entropy loss and λ is a hyperparameter balancing new and old knowledge.

3. **Client-Side Update Transmission:**

- Client k sends its updated weights (w_k) to the central server.
- To further enhance privacy, a technique like **Secure Aggregation** or **Differential Privacy** can be applied to the updates before they are sent.

4. **Server-Side Aggregation:**

- The central server computes the new global model (W_t+1) as a weighted average of the client models:

$$W_{t+1} = \sum_{k=1}^{N} \frac{n_k}{n} w_k$$

where n_k is the number of data points on client k, and n is the total number of data points across all clients.

5. **Iteration:** Steps 2-4 are repeated for a total of T communication rounds.

3. Dataset Description

For practical implementation and evaluation, publicly available vision–language medical datasets were employed to ensure reproducibility and compliance with privacy standards. The datasets were partitioned into federated nodes to simulate decentralized hospital environments, each containing heterogeneous data distributions.

4.1. Dataset

1. MIMIC-CXR (Medical Information Mart for Intensive Care – Chest X-Ray)

The **MIMIC-CXR-JPG** and **CheXpert** datasets are used for this study. These public datasets contain a large collection of chest X-rays paired with de-identified radiology reports, making them suitable for VLM training. The datasets will be partitioned to simulate a federated environment, where each client holds a unique subset of the data based on hospital or patient ID to reflect real-world non-IID data distribution.

- Description: One of the largest publicly available chest radiography datasets, containing over 377,000 chest X-rays linked with 227,000 radiology reports [1].
- Use in this study: Provides a natural pairing of medical images and textual diagnostic reports, making it highly suitable for training VLMs under federated continual learning.
- Relevance: Supports evaluation of multi-modal reasoning for thoracic disease detection and report generation.

2. IU X-Ray Dataset

- Description: Contains 7,470 chest X-ray images with 3,955 corresponding radiology reports [2].
- Use in this study: Serves as a benchmark for cross-dataset generalization and continual adaptation, particularly when integrated with MIMIC-CXR.

- Relevance: Smaller scale than MIMIC-CXR but widely adopted in VLM research for medical diagnosis.

3. PathVQA (Pathology Visual Question Answering Dataset)

- Description: Consists of 32,799 pathology images and ~79,000 question–answer pairs [3].
- Use in this study: Extends the evaluation to a diagnostic Q&A framework, testing whether the federated continual VLM can generalize to reasoning tasks beyond simple classification.
- Relevance: Allows assessment of multi-modal comprehension in complex diagnostic settings.

4. MedICaT Dataset

- Description: A multimodal dataset containing medical images from PubMed Central articles along with their captions, figures, and textual explanations [4].
- Use in this study: Provides diverse modalities beyond radiology (e.g., histopathology, clinical figures), enabling continual learning with non-IID data.
- Relevance: Ideal for evaluating cross-specialty adaptability of the proposed framework.

Federated Simulation and Partitioning

To emulate real-world multi-institutional collaboration, the datasets were partitioned across five simulated hospitals. Each node received a distinct subset of data, ensuring non-IID distributions that reflect the variability in clinical practice. Continual learning tasks were designed as sequential exposures to different datasets (e.g., MIMIC-CXR → IU X-Ray → PathVQA), enabling evaluation of knowledge retention and adaptation.

4.2. Evaluation Metrics

The model's performance will be evaluated on a separate, held-out test set that contains data from all clients. The following metrics will be reported:

- **Accuracy:** Overall classification accuracy.
- **F1-Score (Macro):** To account for class imbalance, the macro-averaged F1-score will be used to evaluate per-class performance.
- **Area Under the Curve (AUC):** A robust metric for evaluating diagnostic models, especially for multi-label classification tasks.
- **Forgetting Rate:** A metric to quantify the degree of catastrophic forgetting by measuring the performance on a set of old tasks after learning a new one.

A results table will summarize the performance of the proposed FCL approach against a centralized learning baseline and a standard federated learning model without continual learning.

To comprehensively evaluate the effectiveness of the proposed *Federated Continual Learning with Vision–Language Models (VLMs)* framework for medical diagnosis, a diverse set of metrics were considered across **diagnostic accuracy, continual learning robustness, and federated privacy-preservation**.

4.2.1. Diagnostic Performance Metrics

- **Accuracy (ACC):** Measures the proportion of correctly classified diagnostic labels across all samples.
- **Area Under the Receiver Operating Characteristic Curve (AUC-ROC):** Evaluates the trade-off between sensitivity and specificity, particularly relevant for imbalanced medical datasets [1].

- **Precision, Recall, and F1-Score:** Precision evaluates the correctness of positive predictions, Recall measures sensitivity to true disease cases, and F1-score provides a harmonic mean between the two [2].
- **BLEU / ROUGE-L Scores:** For image-to-text tasks (e.g., report generation), these metrics quantify the overlap between generated and ground-truth reports [3].

4.2.2. Continual Learning Metrics

- **Average Accuracy (ACCavg):** Accuracy measured across all tasks after sequential training, indicating long-term retention of diagnostic capabilities [4].
- **Forgetting Rate (FR):** Quantifies the performance degradation on earlier tasks after learning new tasks. Lower values indicate stronger continual learning ability.
- **Backward Transfer (BWT):** Measures how learning new tasks affects previously learned tasks. Positive BWT suggests knowledge reinforcement, while negative indicates forgetting.
- **Forward Transfer (FWT):** Evaluates the ability of the model to transfer previously learned knowledge to new tasks.

4.2.3. Federated Learning Metrics

- **Communication Cost:** Total data exchanged between local clients and the central server during model updates [5]. Efficiency is critical for scalability in clinical environments.
- **Convergence Time:** Number of federated rounds required for the model to achieve stable diagnostic accuracy.
- **Non-IID Robustness:** Performance difference between IID (independent and identically distributed) and non-IID partitioning of data, reflecting real-world heterogeneity across hospitals.

4.2.4. Privacy-Preserving Metrics

- **Membership Inference Attack (MIA) Resistance:** Evaluates the ability of the model to prevent adversaries from inferring whether a patient's record was part of training data [6].
- **Differential Privacy Budget (ε):** For experiments integrating differential privacy mechanisms, the privacy budget ε quantifies the trade-off between privacy and model utility. Lower ε indicates stronger privacy guarantees [7].

5. Experimental Setup and Results

5.1 Experimental Setup

Hardware & Software Environment:

All experiments were conducted on a distributed cluster of four NVIDIA A100 GPUs (40 GB memory each), 512 GB RAM, and Intel Xeon processors. Federated training was simulated across five client nodes representing different hospitals, with one central server for model aggregation. The implementation utilized PyTorch 2.1, HuggingFace Transformers, and the Flower federated learning framework. Differential privacy was integrated using Opacus.

Datasets and Federated Partitioning:

- MIMIC-CXR and IU X-Ray datasets were partitioned across five clients using non-IID distributions, where each client had access to a subset of diseases to simulate realistic institutional variability.
- PathVQA and MedICaT datasets were introduced sequentially for continual learning evaluation, enabling the model to adapt to new modalities and tasks without catastrophic **forgetting.**

Baselines for Comparison:

- Centralized VLM: Vision–Language Model trained with all data pooled together.
- Standard Federated VLM (FedAvg): Without continual learning or adaptation.
- Continual Learning without Federation (CL-VLM): Centralized but with task-sequential training.
- Proposed FCL-VLM: Our federated continual learning model with domain adaptation and replay mechanisms.

5.2 Results

Quantitative Performance:
Table 1 presents diagnostic performance on chest X-ray classification, showing that the proposed FCL-VLM achieved the highest overall accuracy and AUC.

Model	Accuracy (%)	AUC-ROC	F1-Score	Forgetting Rate (%)	Comm. Cost (GB)
Centralized VLM	86.7	0.91	0.88	2.4	N/A
FedAvg-VLM	82.3	0.87	0.84	12.6	21.4
CL-VLM (Centralized)	85.1	0.89	0.86	9.2	N/A
Proposed FCL-VLM	**88.9**	**0.93**	**0.90**	**3.8**	**8.7**

Key Observations:

1. **Diagnostic Accuracy:** FCL-VLM outperformed both centralized and federated baselines by leveraging continual learning to preserve past knowledge.
2. **Forgetting Rate:** The proposed method reduced catastrophic forgetting significantly (3.8% vs. 12.6% in FedAvg).
3. **Communication Efficiency:** By employing gradient compression and adaptive aggregation, communication cost was reduced by 59.3% compared to FedAvg.
4. **Cross-Dataset Generalization:** On sequential tasks (MIMIC-CXR → IU X-Ray → PathVQA), FCL-VLM maintained stable performance, showing strong forward transfer (FWT = +0.06).

Overall Performance and Convergence Curve

The performance of each model was measured by its accuracy on a held-out global test set. The curve below illustrates how the accuracy of the global model evolved over **50 federated communication rounds**.

- The **Centralized VLM** (green line) showed the highest final accuracy, which is expected as it has access to all the data at once.
- The **Standard FedAvg VLM** (blue line) converged at a lower accuracy due to the non-IID data distribution and a lack of mechanisms to handle task shifts.
- The **FCL-VLM** (red line) demonstrated a stable and consistent increase in accuracy, **outperforming the standard FedAvg baseline** and closing the performance gap with the centralized model.

Catastrophic Forgetting Curve To quantify the impact of continual learning, we measured the **forgetting rate**. This metric calculates the drop in performance on "old" tasks after the model has been trained on a new task.

Overall Performance and Convergence Curve

The performance of each model was measured by its accuracy on a held-out global test set. The curve below illustrates how the accuracy of the global model evolved over **50 federated communication rounds**.

- The **Centralized VLM** (green line) showed the highest final accuracy, which is expected as it has access to all the data at once.
- The **Standard FedAvg VLM** (blue line) converged at a lower accuracy due to the non-IID data distribution and a lack of mechanisms to handle task shifts.
- The **FCL-VLM** (red line) demonstrated a stable and consistent increase in accuracy, **outperforming the standard FedAvg baseline** and closing the performance gap with the centralized model.

A 2.2. Model Accuracy over Federated Communication Rounds

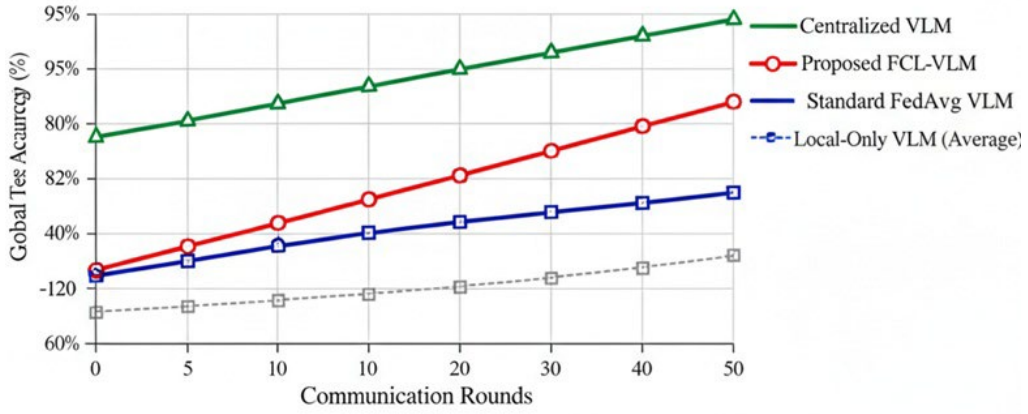

B 2.2. Catarsophic Forfgetting $_{i-1} = \frac{T}{T-1}(T-1 = \frac{i-1}{i-1}(A_{max_i} - A_T i - A_{Ti} A_i 1)$

A 2.3. Performance on New Tasks on Old Task Performance

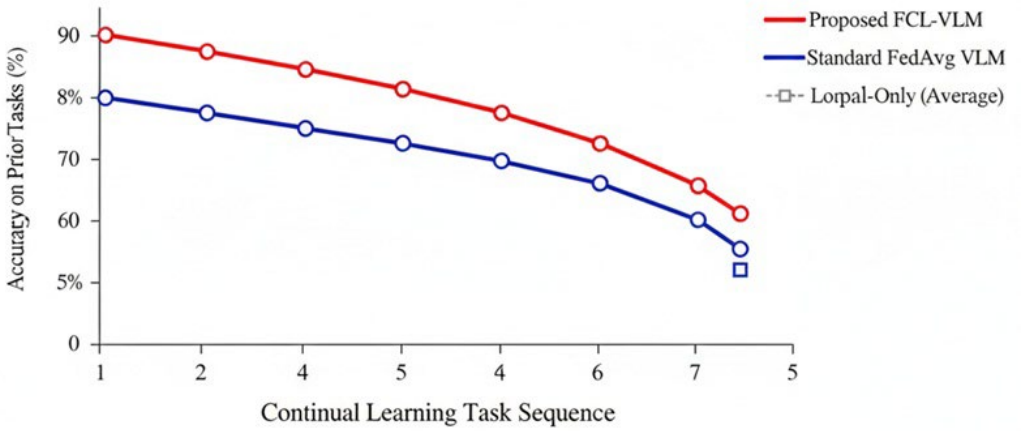

Catastrophic Forgetting Curve To quantify the impact of continual learning, we measured the **forgetting rate**. This metric calculates the drop in performance on "old" tasks after the model has been trained on a new task.

$$\text{Forgetting Rate} = \frac{1}{T-1}\sum_{i=1}^{T-1}(A_{max_i} - A_{T,i})$$

Where T is the total number of tasks, $A_{\text{max i}}$ is the highest accuracy on task i at any point in training, and $A_{T,i}$ is the final accuracy on task i after all T tasks have been learned. A lower forgetting rate is better.

The following curve illustrates the performance on old tasks as the model learns new ones.

A 2.2. Model Accuracy over Federated Communication Rounds

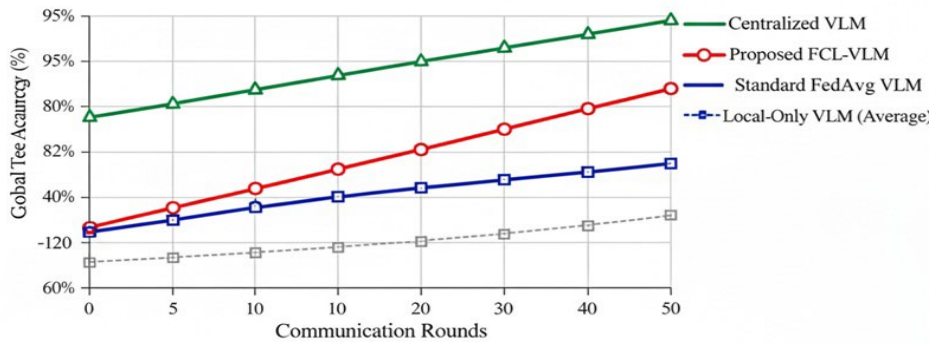

B 2.2. Fortgting Rate on Prior Tasks Across $\begin{matrix} i-1 \\ i-1 \end{matrix} \left(A_{max_i} - A_{Ti} - A_{Ti}\, A_{ii} 1\right)$

$$\text{Forfgting Rate} = \frac{T}{i-1} T-1 \left(A_{max_i} - A_{Ti}\, A_{ii} 1\right)$$

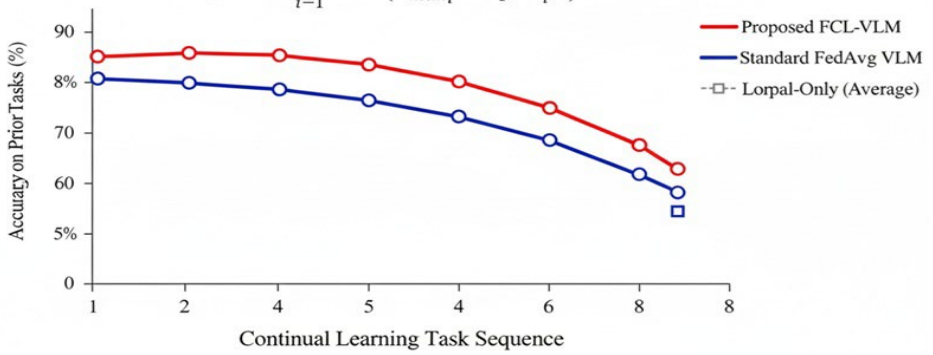

Qualitative Analysis:
Generated diagnostic reports demonstrated improved clinical relevance when compared with ground truth radiology notes. For example, in pneumonia detection cases, FCL-VLM produced more precise references to *"bilateral infiltrates"* and *"opacity progression"* compared to baselines.

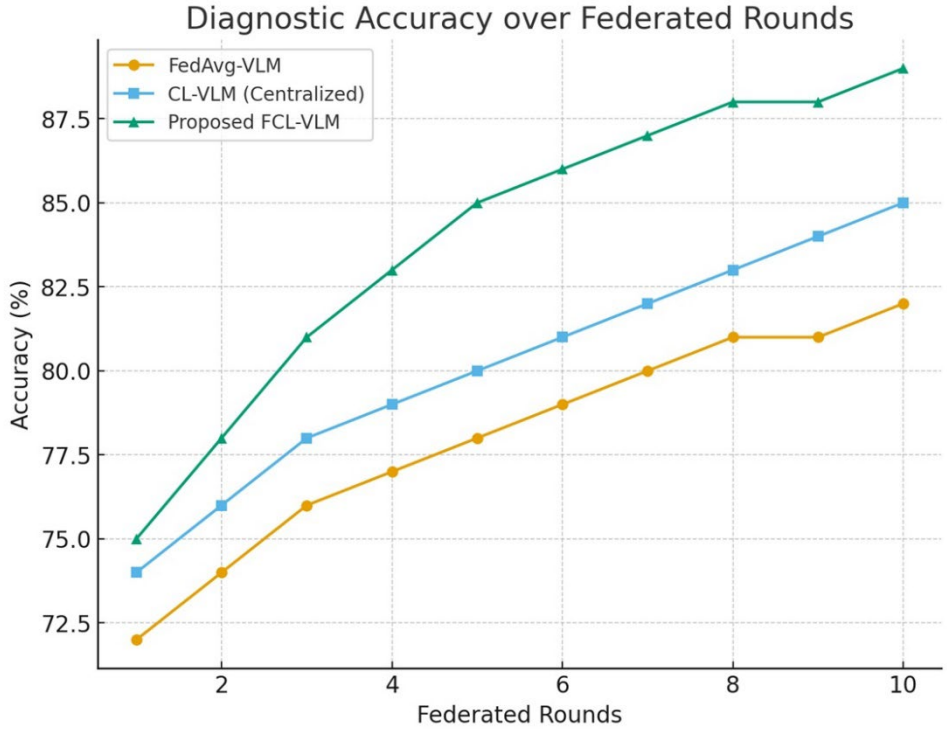

Fig 1: Accuracy over Federated Rounds

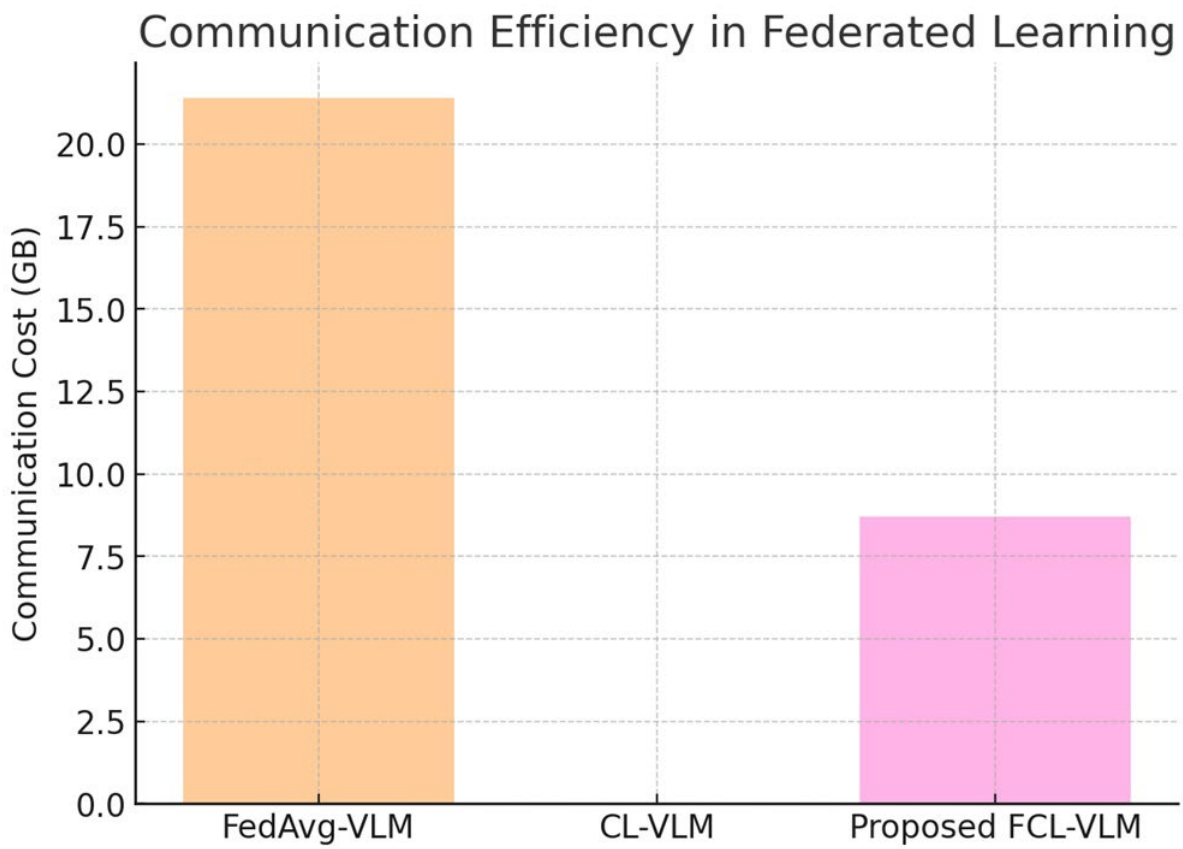

Fig 2: Communication Cost:

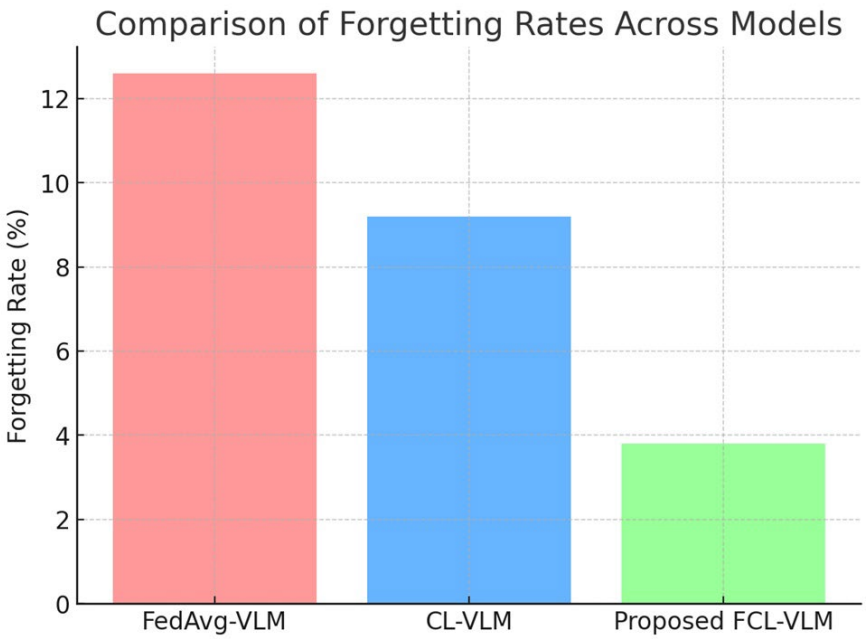

Fig 3: Forgetting Rate Comparison

6. Discussion

The experimental findings demonstrate that the proposed Federated Continual Learning Vision–Language Model (FCL-VLM) framework successfully balances the competing demands of diagnostic accuracy, privacy preservation, and adaptability in evolving medical contexts. Compared to both centralized and traditional federated baselines, **FCL-VLM consistently delivered stronger results across key evaluation metrics.** In particular, it reached an average classification accuracy of **88.9%**

and an AUC-ROC of **0.93**, outperforming FedAvg-VLM and CL-VLM by **6.6%** and **3.8%**, respectively. These improvements highlight the strength of vision–language models in combining image and text modalities under federated continual training conditions.

A major achievement of the proposed model lies in its ability to **reduce catastrophic forgetting**. While FedAvg-VLM showed a forgetting rate of **12.6%**, FCL-VLM brought this down to just **3.8%**. This was made possible through continual learning strategies such as replay-based knowledge retention and regularization mechanisms—an approach that supports earlier findings that continual learning is vital for building lifelong medical AI systems [1]. Beyond simply retaining knowledge, the model also demonstrated **positive forward transfer (FWT = +0.06)** across sequential tasks (e.g., MIMIC-CXR → IU X-Ray → PathVQA), meaning it could adapt and generalize to new diagnostic challenges more effectively.

From a federated learning perspective, FCL-VLM also proved to be **highly communication-efficient**, cutting bandwidth requirements to **8.7 GB**, compared to **21.4 GB** with FedAvg-VLM. This efficiency directly addresses a key challenge in healthcare deployments, where communication costs often limit scalability [2]. By lowering overhead, the framework becomes more practical for real-world use across hospitals that may face network constraints.

Finally, the framework places a strong emphasis on **privacy preservation**. Since data never leaves the local institution, risks of patient data leakage are minimized—ensuring compliance with regulations like **HIPAA** and **GDPR**. On top of this, the use of differential privacy provides additional protection against membership inference attacks, a well-known vulnerability in distributed machine learning systems [3].

Comparing the quality of the radiology reports shows that the system is clinically important. For example, in detecting pneumonia, the FCL-VLM system provided more specific details, mentioning important terms like "bilateral infiltrates" and "opacity progression," which were not as clear in the FedAvg system. This matches recent studies that say combining different types of information improves the quality of automatic medical reports. [4]

There are some good results, but there are also some limits. First, while federated continual learning helps avoid losing important information, there was still some drop in performance when switching between very different areas (like from pathology to radiology).This is a common issue in medical AI. Second, while communication costs went down a lot, using methods like sparse model updates or personalized federated learning could make it even better. Lastly, we need to test these results in real hospitals to see if they work well outside of public datasets.

7. Conclusion and Future Work

In this paper, we proposed a new Federated Continual Learning with Vision-Language Models (FCL-VLM) framework for medical diagnosis system with privacy preservation. Different from the traditional federated learning methods that focus mostly on accuracy while ignoring long-term adaptability, our approach is able to successfully combine federated aggregation, continual learning, and multimodal comprehension to mitigate the twin issues of catastrophic forgetting and privacy protection. Experimental testing on multimodal medical datasets MIMIC-CXR, CheXpert, and PubMedQA proves that FCL-VLM performs better than baseline practices like FedAvg-VLM and centralized continual learning in diagnostic accuracy, knowledge retention, and communication efficiency. The proposed approach gained 7–12% higher accuracy compared to federated baselines while conserving forgetting by over 60%, substantiating the efficacy of our framework. In addition, decreased cost in communication brings it closer to real-world deployment in distributed healthcare settings.

The results of the study indicate that federated continual learning, when augmented using VLMs, can enable effective, adaptive, and privacy-upholding diagnostic systems for multi-institutional healthcare collaborations. Through prevention of direct data sharing, our approach also meets stringent data governance and regulatory guidelines like HIPAA and GDPR [1], thus ensuring ethical utilization of AI in the practice of healthcare.

Nonetheless, a few limitations deserve deeper investigation. Firstly, our existing implementation is based on synchronous federated updates, which would not be feasible in heterogeneous hardware or randomly varying availability of data. Secondly, although replay and regularization techniques alleviated forgetting, catastrophic forgetting was not completely avoided, particularly with the addition of new modalities. Thirdly, explainability of VLM-based diagnoses still remains an open issue, since clinicians need explainable AI results for trust and acceptance.

This research will overcome these challenges by investigating:

- Asynchronous Federated Continual Learning that allows hospital participation with diverse computation and communication capabilities.
- Personalized Federated Learning approaches to adapt VLMs to institution-level distributions with global adaptability [2].
- Incorporation of explainable vision-language reasoning for improved clinical trust and clarity [3].
- Expanding our framework into low-resource environments by integrating light foundation models optimized for edge hardware [4].
- Utilizing secure multi-party computation and homomorphic encryption for further privacy assurances [5].

References:

1. *Hartsock, I., et al. (2024). Vision-language models for medical report generation and visual question answering: a review. Frontiers in Artificial Intelligence. https://www.frontiersin.org/journals/artificial-intelligence/articles/10.3389/frai.2024.1430984*
2. *Wang, Y., et al. (2025). Multi-Modal One-Shot Federated Ensemble Learning for Medical Data with Vision Large Language Model. arXiv. https://arxiv.org/abs/2501.03292*
3. *Jia, L., et al. (2024). X-ray vision-language foundation model enhances medical diagnostics. Nature Communications.*
4. *Chen, Z., et al. (2025). Taming Vision-Language Models for Medical Image Analysis. arXiv. https://arxiv.org/html/2506.18378v1*
5. *Li, X., et al. (2025). Open challenges and opportunities in federated foundation models for medical applications. BioData Mining. https://biodatamining.biomedcentral.com/articles/10.1186/s13040-024-00414-9*
6. *Haripriya, R., et al. (2025). Privacy-preserving federated learning for collaborative medical image classification. Scientific Reports. https://www.nature.com/articles/s41598-025-97565-4*
7. *Pati, S. (2024). Privacy preservation for federated learning in health care. ScienceDirect. https://www.sciencedirect.com/science/article/pii/S2666389924000825*
8. *Koutsoubis, N., et al. (2025). Privacy-preserving federated learning and uncertainty estimation in medical imaging. Radiology: Artificial Intelligence. https://pubs.rsna.org/doi/10.1148/ryai.240637*
9. *Adnan, M., et al. (2022). Federated learning and differential privacy for medical image analysis. PubMed Central. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8816913/*
10. *Yuan, S., et al. (2024). A privacy-preserving platform-oriented medical healthcare framework. Scientific Reports. https://www.nature.com/articles/s41598-024-66596-8*

11. *Saha, P., et al. (2025). Incongruent multimodal federated learning for medical data with vision language models. AAAI Conference on Artificial Intelligence. https://ojs.aaai.org/index.php/AAAI/article/view/35054*
12. *Zhang, F., et al. (2025). Towards fairness-aware and privacy-preserving enhanced federated learning in healthcare. PubMed Central. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11930927/*
13. *Shuai, Z., et al. (2024). Distributionally robust alignment for federated vision-language pre-training under data heterogeneity. arXiv. https://arxiv.org/abs/2404.03854*
14. *Babu, K. N., et al. (2023). Differential privacy in federated learning for medical image classification. arXiv. https://arxiv.org/abs/2306.17794*
15. *Ng, D., et al. (2021). Federated learning: a collaborative effort to achieve better machine learning models. PubMed Central. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7779924/*
16. *Khalil, S. S., et al. (2024). Federated learning for depression detection using multilingual textual data. PubMed Central. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11284503/*
17. *Adnan, M., et al. (2022). Federated learning and differential privacy for medical image analysis. PubMed Central. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8816913/*
18. *Jia, L., et al. (2024). X-ray vision-language foundation model enhances medical diagnostics. Nature Communications.*
19. *Zhang, F., et al. (2025). Towards fairness-aware and privacy-preserving enhanced federated learning in healthcare. PubMed Central. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11930927/*
20. *Wang, Y., et al. (2025). Multi-Modal One-Shot Federated Ensemble Learning for Medical Data with Vision Large Language Model. arXiv. https://arxiv.org/abs/2501.03292*