



RESEARCH ARTICLE

OPEN ACCESS

A NOVEL PRIVACY-PRESERVING DECENTRALIZED FEDERATED LEARNING APPROACH FOR INCREMENTAL HEALTHCARE DATA

Rostom Mennour¹

¹University of Constantine 2 – Abdelhamid Mehri, Faculty of NTIC, LISIA Laboratory, Constantine, Algeria.

¹<https://orcid.org/0000-0001-7317-3672>

Email: rostom.mennour@univ-constantine2.dz

ARTICLE INFO

Article History

Received: May 27, 2025

Revised: June 20, 2025

Accepted: September 30, 2025

Published: October 31, 2025

Keywords:

Federated learning,
Decentralized learning,
Data streams,
Healthcare,
Data privacy.

ABSTRACT

Healthcare institutions face significant challenges in implementing machine learning solutions, particularly in environments with continuous data streams and privacy constraints. Traditional approaches struggle to balance effective model training with data privacy, especially when dealing with varying data volumes across different institutions and the need for continuous adaptation to new medical knowledge. This paper presents a novel decentralized federated continual learning system that enables privacy-preserving collaboration among healthcare institutions without central coordination. Our approach combines Generative Adversarial Networks (GANs) and Deep Stacking Networks (DSNs) in a fully connected network topology, where each node employs a GAN for synthetic data generation and a DSN for classification tasks. The system processes streaming data while maintaining privacy through federated parameter sharing, allowing institutions to benefit from collective knowledge without exposing sensitive patient data. Experimental validation on the MIMIC-VI-ED dataset demonstrates that our approach successfully addresses data volume disparities between institutions, enabling smaller healthcare centers to achieve performance levels comparable to larger institutions. The system demonstrates robust performance comparable to state-of-the-art centralized approaches while providing crucial advantages in terms of data privacy preservation, institutional collaboration, and dynamic data processing capabilities.



Copyright ©2025 by authors and Galileo Institute of Technology and Education of the Amazon (ITEGAM). This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

I. INTRODUCTION

Machine learning holds immense potential to transform healthcare by improving patient care, diagnosis, treatment planning, and disease prevention. By analyzing diverse medical data—such as patient records, diagnostic images, genomic sequences, and real-time monitoring data—machine learning algorithms can uncover patterns and correlations that are often beyond human detection. This capability enables more accurate diagnoses [1, 2], personalized treatments, and early disease detection [3]. Additionally, machine learning enhances operational efficiency in healthcare by optimizing resource allocation, automating administrative tasks, and improving the quality-of-care delivery. The integration of these technologies into healthcare systems promises significant advancements in patient outcomes, cost reduction, and medical research innovation [4, 5].

Models trained on data from a single healthcare system often face generalizability challenges due to variations in demographics, laboratory equipment, EHR systems, data collection frequencies, and clinical protocols. While combining large-scale clinical datasets with deep learning's adaptability offers potential to address these disparities, healthcare data remains fragmented, and issues of accessibility and privacy persist [6]. Machine learning applications require extensive and diverse datasets, which are costly and time-consuming to acquire. Collaboration between hospitals across regions or countries could improve model robustness by incorporating diverse patient profiles and pathologies. However, centralized learning approaches—where data is aggregated into a single repository—face operational challenges such as large-scale data collection, resource sharing, and heightened privacy risks. Protecting patient data is

critical, especially given the increasing frequency of healthcare data breaches. These challenges highlight the need for innovative frameworks that balance collaboration with privacy preservation.

Federated learning (FL) [7] is a transformative approach to machine learning that enables collaborative model training across institutions while preserving data privacy. In its conventional form, FL relies on a centralized server to coordinate the process. However, centralized FL has notable limitations. It introduces a single point of failure, making the entire system vulnerable to disruptions if the server malfunctions. Additionally, it creates trust dependencies, as participants must rely on the server to maintain data privacy and integrity. The central server also often becomes a bottleneck when managing updates from numerous clients, leading to delays and computational inefficiencies [8]. Furthermore, establishing a central server is complicated by legal and regulatory constraints, especially when data from multiple institutions is involved [9].

Moreover, healthcare data is expanding rapidly due to the widespread adoption of sensing devices and data collection tools, generating vast, real-time data streams globally [10]. Traditional FL methods often rely on rigid assumptions, requiring predefined and static data and class structures, which conflict with the dynamic nature of real-world healthcare [11]. Extracting insights from streaming processes poses challenges due to their dynamic nature, which significantly impacts the entire analysis workflow [12], and training separate models for each task—demands excessive computational resources and fails to leverage shared representations effectively [13]. This dynamic nature of healthcare data necessitates an incremental learning approach that allows models to update continuously without retraining from scratch [14].

Addressing these challenges requires a machine learning solution that enables collaboration among clinical institutions while ensuring data privacy and adapting to new knowledge without losing previously acquired insights. To this end, we propose a decentralized federated learning system where hospital nodes collaboratively train models without relying on a central server. Each node independently operates generative and discriminative models, leveraging streaming data to generate synthetic data and improve performance. Nodes periodically exchange generative model parameters, integrating these updates into their local models to benefit from the collective knowledge of all participating nodes. This paper is organized as follows: Section 2 reviews related work in these domains within healthcare contexts. Section 3 details the proposed decentralized federated continual learning approach. Section 4 presents experimental results validating the system's effectiveness using the MIMIC-IV-ED dataset. Finally, Section 5 concludes with key findings and future research directions.

II. RELATED WORKS

This section reviews the existing literature on federated learning (FL) and continual learning (CL) within the context of healthcare. Federated learning has emerged as a critical approach to addressing privacy concerns and data silos in healthcare by enabling collaborative model training across multiple institutions without sharing sensitive patient data. Concurrently, continual learning offers solutions to the challenges posed by the dynamic and evolving nature of medical data, allowing models to incrementally learn from new data without forgetting previously acquired knowledge. We will explore seminal and recent contributions in these fields, highlighting their applications, benefits, and limitations in healthcare settings.

According to [15], all the articles about the application of federated learning in the healthcare sector were published after 2018, showing an exponential increase in the number of relevant publications over the past few years. Specifically, there was 1 article (0.2% of the total) in 2018, and 253 articles (41.3%) in the first three quarters of 2023. Among the three forms of data partitioning in federated learning, horizontal data partitioning was the most prevalent. Additionally, a significant majority of the articles (83.7%) used a centralized communication architecture, while only 10.6% adopted a decentralized approach. We begin by reviewing several key papers that explore the application of federated learning in healthcare, highlighting how this approach has been utilized to address privacy concerns and improve collaborative medical research.

In [16], the authors applied a federated learning technique for the diagnosis of breast cancer. they present a memory-aware curriculum learning method for federated learning, designed to improve local model consistency by penalizing forgotten samples. This approach prioritizes training samples that are forgotten after the global model's deployment and incorporates unsupervised domain adaptation to handle domain shifts while preserving data privacy. The study in [17] proposes FedMRI, a specificity-preserving federated learning algorithm for MR image reconstruction. The core concept is to split the MR reconstruction model into two parts: a globally shared encoder that generates a generalized representation, and a client-specific decoder that maintains each client's unique domain-specific properties. The researchers introduce in [18] the first federated learning framework for multiple sclerosis (MS) lesion segmentation, utilizing two effective re-weighting mechanisms. The framework assigns a learnable weight to each local node during aggregation based on its segmentation performance.

Additionally, the segmentation loss function for each client is re-weighted according to the lesion volume in the training data. Kassem et al. introduce FedCy [19], a federated semi-supervised learning method that integrates federated learning with self-supervised learning to enhance surgical phase recognition using both labelled and unlabelled videos in a decentralized dataset. FedCy leverages temporal patterns in the labelled data to guide the unsupervised training on unlabelled data. The largest federated learning study to date was presented in [20], encompassing data from 71 sites across 6 continents, to develop an automatic tumour boundary detector for glioblastoma, a rare disease. In [21], the authors explore the potential of edge computing in medicine by evaluating intelligent clinical data processing at the edge. they implemented clustered federated learning (CFL) to develop an automatic COVID-19 diagnosis system.

Several other studies have utilized federated learning for various healthcare applications, including predicting mortality in intensive care units [22], analyzing electronic health records [23], and classifying tumor-infiltrating lymphocytes [24]. All the mentioned articles have employed centralized federated learning, a technique associated with several limitations as outlined in the introduction. However, only a few studies have explored a distributed configuration of federated learning. FedDIS [25], a novel federated learning approach for medical image classification, aims to reduce non-IIDness across clients by generating data locally at each client while sharing medical image data distribution from others, all while safeguarding patient privacy. In their paper [26], the researchers introduce a peer-

to-peer federated learning (P2PFL) framework utilizing Vision Transformers (ViT) models. Their aim is to address key challenges and accurately classify COVID-19 versus normal cases in Chest X-ray (CXR) images. The study in [27] introduces a decentralized federated learning algorithm based on deep neural networks, aiming to mitigate concerns regarding untrusted central servers.

The approach decomposes the original problem into subproblems with consensus constraints, solvable through local computation and communication. The authors of [28] introduce a fully decentralized federated learning method employing knowledge distillation to uphold data privacy and security. Each node operates autonomously without requiring access to external data sources. DEEP-FEL, a decentralized, efficient, and privacy-enhanced federated edge learning system, is introduced in [29]. DEEP-FEL allows medical devices across different institutions to collaboratively train a global model without exchanging raw data. It employs a hierarchical ring topology to reduce centralization and formulate the ring construction as an optimization problem, solvable through an efficient heuristic algorithm. While some studies have adopted federated learning approaches without relying on a central server, none of the above cited articles have addressed the dynamic nature of data in the healthcare field.

However, the exploration of continuous learning has begun to tackle this aspect, highlighting the need for models that can adapt to the ever-evolving medical data landscape. The survey in [30] aims to provide an overview of the current techniques, applications and challenges, of continual learning in physiological signal analysis. It identified eight articles focusing on the application of continual learning methods to physiological signals published between January 2021 and September 2023. Deep Generative Feature Replay (DGFR), a continual learning approach for cancer classification, is proposed in [31]. DGFR includes an Incremental Feature Selection (IFS) module to identify the most significant CpG sites from high-dimensional data, optimizing the number of selected sites.

It also features a Scholar Network (SN) that employs a variational autoencoder (VAE) to generate pseudo data without needing past samples and a neural network classifier to predict cancer types. The study in [32] introduces Lifelong nnU-Net, a standardized framework designed to facilitate continual segmentation for researchers and clinicians. Based on the highly regarded nnU-Net, this framework includes all necessary modules for sequential training and testing, ensuring broad applicability and making it easier to evaluate new methods in a continual learning context. The review in [33] offers a comprehensive overview of state-of-the-art continual learning techniques in medical image analysis. It covers key issues such as catastrophic forgetting, data drifts, and the balance between stability and plasticity. It discusses continual learning frameworks, including scenarios, techniques, evaluation schemes, and metrics. The review also addresses challenges like costly data annotation, temporal drift, and the need for benchmarking datasets.

Despite the increasing popularity of federated learning and continuous learning techniques in healthcare, the intersection of these two fields remains largely unexplored, with only a few articles addressing their combined potential. The study in [34] explores the feasibility of continual learning for multicenter collaboration in brain metastasis identification. The synaptic intelligence (SI) continual learning algorithm is employed to preserve important model weights while sequentially training across different centers. Sequential training faces several limitations, including significant communication overhead and scalability issues as the number of participating nodes increases. Synchronization challenges arise due to heterogeneous environments, and fault tolerance is a concern since the failure of a single node can disrupt the entire process. A privacy-preserving algorithm called WUPERR (Weight Uncertainty Propagation and Episodic Representation Replay) for early sepsis prediction is presented in [4]. Continual learning is used for federated learning purpose, but the data in this configuration is considered to be static.

III. PROPOSED APPROACH

In this article, we propose a novel approach for decentralized continual federated learning for emergency department in healthcare institutions.

III.1 DATA DESCRIPTION AND PREPROCESSING

To train and evaluate our federated continual architecture, we utilized the MIMIC-VI-ED dataset [35]. MIMIC-IV-ED is a comprehensive, publicly accessible database of approximately 425,000 emergency department admissions at Beth Israel Deaconess Medical Center from 2011 to 2019. It includes vital signs, triage information, medication details, and discharge diagnoses, all deidentified in compliance with HIPAA regulations. We then performed a preprocessing phase according to the method outlined in [36]. The preprocessed dataset was then horizontally divided into subgroups, with each subgroup representing the data for a single healthcare institution. During the training phase, the algorithm for each institution does not access the entire dataset at once but receives it progressively in a streaming manner.

III.2 PROBLEM FORMULATION

In this paper, we address the problem as follows: Consider N hospitals (nodes) in a decentralized, fully connected network topology, without a central server. Each hospital is labeled by $i \in \{1, 2, \dots, N\}$. The dataset at hospital i is represented as $D_i = \bigcup_{t=1}^{\infty} D_i^t$, where t indicates the time when the data was received. Here, $D_i^t \cap D_i^r = \emptyset$ if $t \neq r$ and $D_i \cap D_j = \emptyset$ if $i \neq j$. Each hospital can only access its own dataset and does not have access to the data of other hospitals. At any given timestamp t , a hospital can only see the data D_i^t and cannot access historical or future data. The objective is to effectively train N continual learning models on their respective private task streams by communicating model parameters among themselves in a fully distributed manner.

III.3 GLOBAL SYSTEM ARCHITECTURE

In this decentralized federated learning system, multiple hospitals act as nodes in a fully connected peer-to-peer network without a central server. Each hospital independently receives and processes streaming data in real-time, performing continual learning on the incoming data. As illustrated in Fig. 1, the architecture ensures that each hospital node is directly linked to every other node in the network.

Hospitals periodically train their local models using the latest data collected, ensuring that models remain up-to-date with current information. After a set number of local updates, each hospital shares its updated model parameters with all other nodes in the network. Upon receiving these parameters, each hospital aggregates them using the method detailed in a later section, integrating knowledge from multiple sources to enhance its local model with insights from other hospitals' data. This iterative cycle of local training, model sharing, aggregation, and updating enables continuous adaptation and improvement over time. Importantly, patient data remains local to each hospital, ensuring privacy and compliance with regulations, as only model parameters—not raw data—are exchanged. This decentralized approach leverages the collective intelligence of all participating nodes while safeguarding sensitive patient information, making it a robust and privacy-preserving solution for collaborative healthcare analytics and machine learning.




Figure 1: Global system architecture.
Source: Authors, (2025).

III.3 NODE LEVEL PROCESS

Each node in the decentralized federated learning system operates with several key components: a generative model to recreate data from previous streams, a discriminative model to classify data, and a vector to store the generative model parameters received from other nodes. Periodically, each node receives a new batch of data. At this point, the node uses its generative model to recreate older data from previous streams and combines it with the newly received batch. This combined dataset is then used to update the generative model, ensuring it can generate accurate synthetic data in future iterations. Additionally, the node uses the generative models stored in its vector—representing data streams seen by other nodes—to generate further synthetic data. This synthetic data is combined with the node's local dataset, after which the discriminative model is retrained on this comprehensive dataset. Finally, the updated generative model is broadcast to all other nodes, which store these parameters in their respective vectors. This process allows each node to continuously improve its models by integrating knowledge from both its own data and synthetic data representing other nodes' experiences. As a result, the system enhances overall performance while maintaining strict data privacy.

III.3.1 CONTINUAL LEARNING ASPECT

Continual learning at each node ensures that models adapt and improve as new data becomes available. Periodically, each node receives new data batches from its local sources. To effectively integrate this data, the node first uses its generative model to recreate historical data from previous streams, generating a comprehensive dataset that combines both past and newly received data. This combined dataset is essential for updating the generative model, ensuring it captures both historical and current data trends, thereby enhancing its ability to simulate future data accurately. The node then retrains its discriminative model using this comprehensive dataset. By incorporating both newly received and generated historical data, the discriminative model continuously adapts to evolving information, improving its classification and prediction capabilities over time. This iterative process, illustrated in Fig. 2, enables each node to stay aligned with changing data patterns, ensuring that models remain accurate, robust, and relevant in dynamic environments.

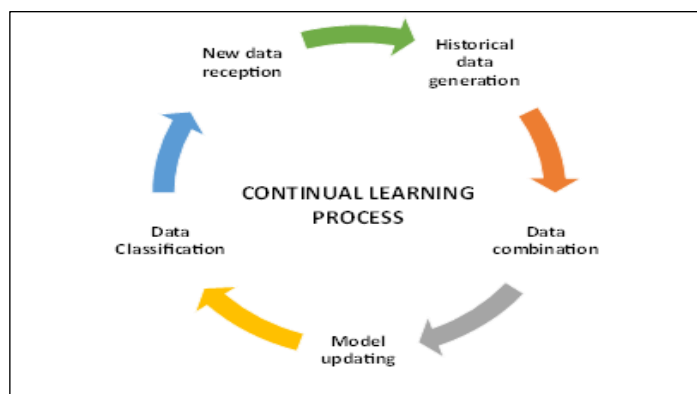


Figure 2: Continual learning process.
Source: Authors, (2025).

III.3.2 FEDERATED LEARNING ASPECT

The federated learning component of the system facilitates collaborative knowledge sharing among multiple nodes without relying on a central server. Each node in the network maintains a vector that stores generative model parameters received from other nodes. Periodically, nodes update their generative models with newly received data and broadcast these updated models to all other nodes in the network. Upon receiving these updates, each node refreshes the corresponding entries in its vector, ensuring it has the latest model parameters from all participating nodes. Using these stored generative models, each node generates synthetic data that represents the data streams observed by other nodes. This synthetic data is then combined with the node's own generated and newly received data to create a diverse and enriched dataset for retraining its discriminative model. By incorporating synthetic data from other nodes, each node effectively simulates and learns from experiences across the network, enhancing its own model's performance. This process, illustrated in Fig. 3, occurs continuously, enabling all nodes to benefit from the collective intelligence of the network while maintaining strict data privacy and security. Through this collaborative approach, the federated learning system leverages diverse insights from all nodes, resulting in improved overall system performance and more generalized models.

III.3.3 NODE'S COMPONENTS

Each node in the proposed system consists of two primary models and a storage vector for the models of other nodes, as illustrated in Fig. 4. The primary models include a Generative Adversarial Network (GAN) [37], which serves as the generative model, and a Deep Stacking Network (DSN) [38], which functions as the discriminative model. Each component is examined individually and described in detail in the following sections.

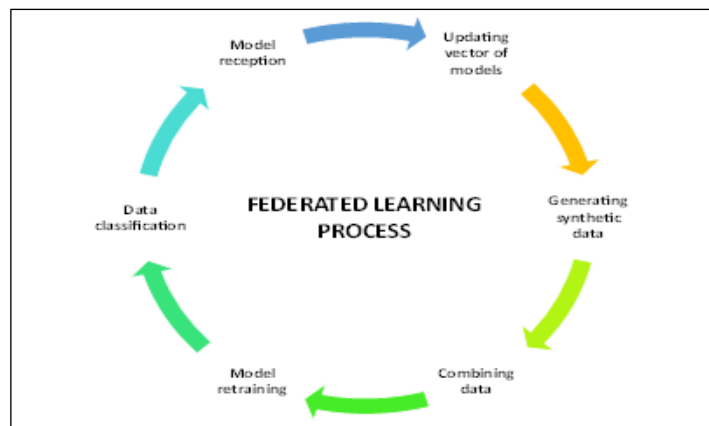


Figure 3: Federated learning aspect.

Source: Authors, (2025).

a) Generative model

In the proposed system, a Generative Adversarial Network (GAN) is employed as the generative model. GANs are a class of neural network architectures designed to generate synthetic data that closely resembles the real data on which they are trained. The GAN comprises two sub-models: the generator and the discriminator. The system is designed to train the GAN on a combination of newly received data and historical data, enabling it to continuously adapt and refine its data generation capabilities by incorporating both current and past information. Newly received data represents the latest batch of information collected by the system, reflecting current trends and patterns, while historical data refers to synthetic data generated by the GAN based on its prior learning from earlier data streams.

Both components of the GAN—the generator and discriminator—are implemented as Multi-Layer Perceptrons (MLPs). During each training iteration, the new data is combined with historical data generated by the GAN itself to form a comprehensive training dataset. The generator is trained to produce synthetic data that not only mirrors the characteristics of the new data but also retains features from historical data. Its objective is to improve its output such that the discriminator cannot differentiate between real data (a combination of new and historical) and synthetic data. Meanwhile, the discriminator is tasked with distinguishing between real and synthetic data, receiving both types during training.

The training process follows an adversarial framework, where the generator and discriminator compete against each other in a feedback loop. As the generator improves its ability to create realistic synthetic data, the discriminator becomes increasingly proficient at identifying fake data. This adversarial dynamic ensures that the GAN continuously enhances its generative capabilities while maintaining alignment with evolving data trends, making it a critical component of the system's continual learning framework.

b) Discriminative model

The discriminative model in the proposed system is implemented as a Deep Stacking Network (DSN), a robust and scalable neural network architecture designed for complex classification tasks. The DSN consists of three identical modules, each implemented as a Multi-Layer Perceptron (MLP) with three layers: an input layer, a hidden layer, and an output layer. The first module receives input that combines local data—including newly received data and historical data generated by the node's own generative model—with external data synthesized by the generative models stored in the vector of models, representing data from other nodes in the federated network. This integration ensures that the model captures both local trends and broader patterns observed across the system. The hidden layer in each module transforms the input data into higher-level features, enabling the extraction of intricate patterns essential for accurate classification. The output layer of each module generates intermediate results, which are passed sequentially to subsequent modules. The

outputs of the first and second modules serve as inputs to the second and third modules, respectively, allowing the DSN to progressively refine its understanding of the data through each stage. The final classification is performed by the third module, ensuring that the model achieves high accuracy by leveraging both local and collaborative knowledge. This hierarchical structure enables the DSN to effectively handle complex classification tasks while maintaining adaptability in dynamic and distributed.

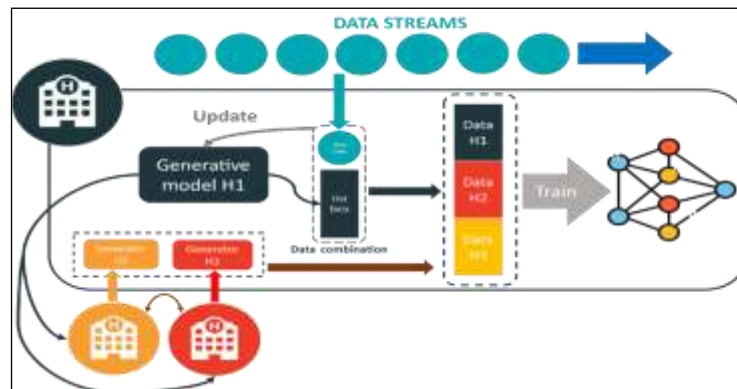


Figure 4: Data path through the different components of a node.
Source: Authors, (2025).

IV. RESULTS AND DISCUSSIONS

This section presents a comprehensive evaluation of the proposed decentralized federated continual learning system through a series of experiments designed to validate its effectiveness and analyze various aspects of its functionality. The evaluation begins with an ablation study to isolate and assess the contributions of individual components. Specifically, we evaluate the performance of the Deep Stacking Network (DSN) in a centralized configuration, examine the streaming process without collaboration in both single and multi-institution settings, and analyze decentralized federated learning without streaming. Following the ablation study, we evaluate the complete system, integrating all components—streaming, and decentralized collaboration.

Beyond validating the system's overall performance, we conduct targeted experiments to explore specific aspects of the approach. These include studying scalability by assessing the effect of adding more institutions on performance, analyzing the impact of synthetic data size on learning efficiency, evaluating knowledge retention to measure how well the system preserves insights from previous data streams, and investigating knowledge diffusion to quantify how effectively learned knowledge propagates across clients. Finally, we compare our approach with several state-of-the-art methods, including FedAvg, FedProx, FedCIL, and Oracle.

IV. 1 CENTRALIZED DSN EVALUATION

In our first experiment, we evaluate the baseline performance of the Deep Stacking Network (DSN) architecture using the complete MIMIC-VI-ED dataset at a single institution. The DSN is configured with three sequential modules, each containing input, hidden, and output layers, where the output of each module serves as input for the subsequent one. We train the model for a total of 3000 epochs across all modules combined, using 80% of the dataset for training while reserving 20% for testing. The performance metrics, expressed in terms of precision and recall as shown in the Fig. 5, establish our reference point for comparing subsequent streaming and federated approaches, demonstrating the model's fundamental classification capabilities when trained on a complete, static dataset before introducing the complexities of streaming data processing and inter-hospital collaboration. The baseline DSN model demonstrates robust performance on the complete dataset, achieving a precision of 0.84 and recall of 0.8, indicating strong classification capabilities in both correctly identifying positive cases and minimizing false negatives.

IV.2 STREAMING PROCESS EVALUATION IN A SINGLE AND MULTI-INSTITUTION SETTING

The second experiment evaluates our system's performance in a streaming environment across two scenarios. In the first scenario, we assess the model's ability to handle sequential data streams at a single institution, where the complete dataset is delivered in 20 consecutive streams. The system processes these streams sequentially using the complete architecture (GAN for synthetic data generation and DSN for classification). It employs continual learning mechanisms to integrate new information while preserving knowledge from previous streams through synthetic data generation.

In the second scenario, we extend the evaluation to multiple institutions by distributing the data across five hospitals with varying data volumes to reflect real-world healthcare settings. The data distribution follows an uneven pattern (30%, 20%, 20%, 15%, 15%) to simulate practical situations where larger institutions, such as major hospitals, have access to more extensive datasets, while smaller health centers possess significantly less data. Each hospital processes its streams independently without any inter-hospital collaboration, allowing us to analyze how data volume disparities affect model performance in a multi-institutional setting. This setup helps understand the system's behavior under realistic healthcare scenarios where data availability varies significantly across institutions.

The results of first scenario demonstrate strong performance, achieving a precision of 0.83 and recall of 0.8. When compared to the baseline DSN experiment, where the model had access to the complete dataset at once, the streaming approach maintains comparable performance levels. This indicates that our system effectively handles sequential data processing without significant degradation in classification capabilities. As shown in the Fig. 6, the model successfully integrates new information from each stream while maintaining its predictive power.

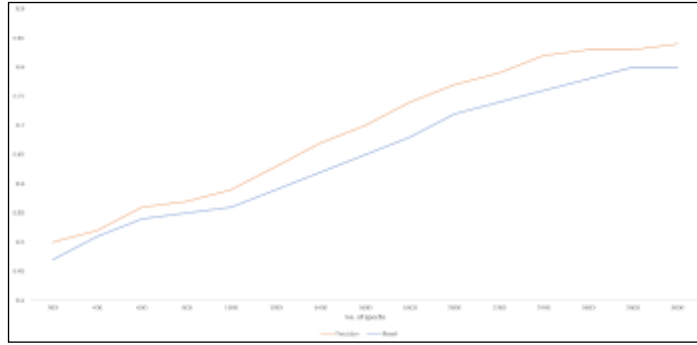


Figure 5: Precision and recall of traditional DSN model on 100% of available data in a single institution configuration. Source: Authors, (2025).




Figure 6: Precision and recall evolution of streaming architecture experiment in a single institution configuration. Source: Authors, (2025).

In the multi-institutional scenario, performance varies significantly across hospitals of different sizes. The largest hospital (30% data volume) achieves the highest performance with a precision of 0.72 and recall of 0.70. Medium-sized hospitals (20% data volume) show moderate performance with a precision of 0.66 and recall of 0.63, while smaller hospitals (15% data volume) demonstrate lower metrics with a precision of 0.58 and recall of 0.56. These results, when compared to the single-institution scenario (precision: 0.83, recall: 0.8), indicate that data volume significantly impacts model performance. For visualization clarity in the Fig. 7 and Fig. 8, we aggregated results from hospitals with similar data volumes, presenting three representative curves: one for the largest hospital (30%), one averaged curve for medium-sized hospitals (20%), and one for smaller hospitals (15%). This representation clearly illustrates how institutional data volume correlates with model performance in a streaming environment without collaboration.




Figure 7: Precision evolution of streaming system architecture in the scenario of five clients without collaboration. Source: Authors, (2025).




Figure 8: Recall evolution of streaming system architecture in the scenario of five clients without collaboration. Source: Authors, (2025).

IV.3 DECENTRALIZED FEDERATED LEARNING EVALUATION WITHOUT STREAMING

In the decentralized federated learning configuration, where five hospitals collaborate without streaming data, the results demonstrate significant performance improvements across all institutions regardless of their data volume. The largest hospital (30% data volume) achieves a precision of 0.83 and recall of 0.79, while medium-sized hospitals (20% data volume) show comparable performance with a precision of 0.83 and recall of 0.78. Notably, smaller hospitals (15% data volume) maintain similar performance levels with a precision of 0.82 and recall of 0.78. These results reveal that federated collaboration effectively neutralizes the performance disparities observed in the non-collaborative scenario, enabling smaller institutions to achieve performance metrics nearly identical to larger ones. As shown in the Fig. 9, where results are aggregated for hospitals with similar data volumes, the convergence patterns demonstrate remarkable consistency across institutions of different sizes, highlighting the effectiveness of our decentralized federated learning approach in equalizing performance across participating hospitals.




Figure 9: Evaluation results of the decentralized federated architecture without streaming in a configuration of five institutions. Source: Authors, (2025).

IV.4 PERFORMANCE ANALYSIS OF THE FULLY INTEGRATED SYSTEM

The final experiment evaluates the complete system incorporating both federated learning and streaming capabilities across five hospitals with varying data volumes. The system demonstrates significant evolution in performance across all institutions throughout the streaming process. Initially, all hospitals show relatively low performance metrics, with the largest hospital (30% data volume) starting at precision 0.42 and recall 0.29, medium-sized hospitals (20% data volume) at precision 0.41 and recall 0.27, and smaller hospitals (15% data volume) at precision 0.38 and recall 0.26. As the system processes more streams and hospitals continue to collaborate, performance steadily improves across all institutions, ultimately converging to comparable high levels regardless of institutional size. By the final stream, the largest hospital achieves precision 0.81 and recall 0.78, while medium-sized and smaller hospitals demonstrate similar performance with precision 0.82/0.81 and recall 0.78/0.77 respectively. As illustrated in the Fig. 10 and Fig. 11, where results are aggregated for hospitals with similar data volumes, the combination of federated learning and streaming capabilities enables all institutions to achieve high performance levels, effectively neutralizing the initial performance disparities related to data volume differences.




Figure 10: Evolution of precision at each institution in the decentralized federated learning system as data streams arrive. Source: Authors, (2025).




Figure 11: Evolution of recall at each institution in the decentralized federated learning system as data streams arrive. Source: Authors, (2025).

IV.5 SCALABILITY: EFFECT OF ADDING MORE INSTITUTIONS

To assess the scalability of our decentralized federated learning system, we conducted experiments by progressively increasing the number of client nodes and evaluating precision and recall after processing 20 data streams. The results indicate a slight decrease in system performance as the number of clients increases like shown in Fig. 12. Specifically, the average precision decreased from 0.84 when using a single institution to 0.82 with five clients, 0.81 with ten clients, and 0.8 with twenty clients. Similarly, the average recall followed a comparable trend. This degradation is primarily attributed to the division of the dataset among clients, resulting in smaller data volumes per client as the number of clients increases. This reduction in local data availability negatively impacts the efficiency of each client's model, leading to a gradual decline in overall system performance. Despite this decrease, the system maintains robust performance across different scales, demonstrating its ability to adapt to various multi-institutional settings while preserving data privacy and facilitating collaborative learning.




Figure 12: System scalability as the number of nodes increases.
Source: Authors, (2025).

IV.6 SYNTHETIC DATA SIZE: IMPACT ON LEARNING EFFICIENCY

In this subsection, we analyze how the size of synthetic data generated by the system influences the efficiency of the learning process. The experiment evaluates synthetic data proportions ranging from 0% to 75% of the original data size. At 0%, no synthetic data is generated, which means that neither collaboration nor incremental processing is possible. Consequently, the system's performance is significantly limited, achieving only a precision of 0.55 and recall of 0.49. As synthetic data generation increases, the system begins to leverage collaborative and incremental learning, resulting in improved performance like illustrated in Fig. 13. The optimal results are observed when synthetic data constitutes between 45% and 60% of the total training data, where precision and recall reach their peak values. Beyond this range, however, performance starts to decline, likely due to over-reliance on synthetic data, which may reduce diversity or introduce noise into the learning process.




Figure 13: Evolution of precision and recall of the system as synthetic data size changes
Source: Authors, (2025).

IV.7 KNOWLEDGE RETENTION: ABILITY TO PRESERVE INSIGHTS FROM PREVIOUS STREAMS

The objective of this experiment is to evaluate the system's ability to retain knowledge from previous data streams while training on newly received data batches. This experiment is motivated by the challenge of catastrophic forgetting, a common issue in continual learning where models lose previously acquired insights when adapting to new information. The experiment is conducted in two phases: first, a pre-update evaluation measures the model's performance on previous streams before training on the new batch; second, a post-update evaluation assesses performance after integrating the new batch, quantifying knowledge retention and identifying potential degradation. Both evaluations (pre and post update evaluations) used a test dataset derived exclusively from previous streams, excluding the newly received data. The experiment was conducted at multiple time points corresponding to the arrival of streams 2, 5, 10, 15, and 20, with tests performed on data generated from older streams. The results, illustrated by Fig. 14 and Fig 15, showed that while there is a slight decline in precision and recall metrics post-update compared to pre-update, the differences were minimal (mean differences of 0.018 for precision and 0.011 for recall), with low variability (standard deviations of 0.0084 for precision and 0.0055 for recall).




Figure 14: Pre and post update precision evolution over time during the knowledge retention experiment. Source: Authors, (2025).




Figure 15: Pre and post update recall evolution over time during the knowledge retention experiment. Source: Authors, (2025).

IV.8 KNOWLEDGE DIFFUSION: PROPAGATION OF KNOWLEDGE ACROSS INSTITUTIONS

This experiment aims to evaluate the extent of knowledge diffusion among clients in the decentralized federated learning system. Specifically, it investigates whether the knowledge learned by a client on its local data propagates effectively to other clients, enabling them to perform well on data streams they have not directly observed. To study knowledge diffusion in the decentralized federated learning system, we evaluate each client's ability to perform on test datasets derived exclusively from other clients' local data. After training, each client's discriminative model is tested on unseen data streams from other clients, and the results are compared against those obtained from exclusive local training (without collaboration).

The evaluation is conducted using average precision and recall as metrics to quantify the effectiveness of knowledge transfer across clients. The results of the knowledge diffusion experiment in Fig. 16 demonstrate the effectiveness of collaborative learning in improving client performance on unseen data. Under collaborative training, the average precision and recall across clients reached 0.82 and 0.79, respectively, compared to 0.77 and 0.74 in the local training scenario without collaboration. This significant improvement highlights the ability of the decentralized federated learning system to propagate knowledge between clients effectively, enabling each node to benefit from insights learned by others.




Figure 16: Comparison of collaborative and local average performances in the knowledge diffusion experiment. Source: Authors, (2025).

IV.9 COMPARATIVE ANALYSIS

In this subsection, we compare the performance of our proposed decentralized federated continual learning approach with several state-of-the-art methods, including FedAvg [7], FedProx [39], FedCIL [40], [41], and Oracle. FedAvg and FedProx represent two widely used aggregation techniques in federated learning, serving as benchmarks for collaborative model training. FedCIL employs GANs to generate samples from previous tasks. Finally, Oracle serves as an upper bound for comparison, assuming clients have access to all their historical data during training.

The results in Table. I. demonstrate that our approach achieves comparable performance to FedCIL, with both methods showing high precision and recall. However, our decentralized framework eliminates the reliance on a central server, addressing critical limitations such as single points of failure and bottlenecks in communication. While Oracle, as an upper bound, outperforms all methods due to its access to all historical data, our method closely approaches its performance while preserving privacy and ensuring scalability. In contrast, FedAvg and FedProx show significantly lower performance, highlighting their limitations in handling continual learning tasks.

Table 1: Comparative results of our approach with FedAvg, FedProx, FedCIL and oracle

Algorithm	Avg Precision	Avg Recall
FedAVG	0.7	0.67
FedProx	0.74	0.72
FedCIL	0.82	0.78
Ours	0.81	0.78
Oracl	0.85	0.81

Source: Authors, (2025).

V. CONCLUSIONS

In this paper, we addressed the critical challenges facing healthcare institutions in implementing machine learning solutions: data privacy concerns, continuous data streams processing, and effective collaboration across institutions of varying sizes. We proposed a novel decentralized federated continual learning system that combines GANs for synthetic data generation and DSNs for classification, enabling privacy-preserving collaboration without central coordination. Experimental results demonstrate the system's effectiveness across diverse scenarios, including ablation studies, scalability analysis, and comparisons with state-of-the-art approaches. The results are comparable with state-of-the-art centralized approaches, and while the system may not significantly improve classification performance beyond existing methods, it provides crucial advantages in terms of data privacy preservation, institutional collaboration, and dynamic data processing capabilities. Furthermore, the analysis of knowledge retention and diffusion highlights the system's ability to preserve insights from historical data and propagate knowledge across institutions, ensuring equitable performance even for smaller healthcare centers. These characteristics make our approach particularly suitable for real-world healthcare applications where privacy, continuous learning, and inter-institutional collaboration are essential requirements.

VI. AUTHOR'S CONTRIBUTION

Conceptualization: Rostom Mennour.
Methodology: Rostom Mennour.
Investigation: Rostom Mennour.
Discussion of results: Rostom Mennour.
Writing – Original Draft: Rostom Mennour.
Writing – Review and Editing: Rostom Mennour.
Resources: Rostom Mennour.
Supervision: Rostom Mennour.
Approval of the final text: Rostom Mennour.

VII. REFERENCES

- [1] E. Aslan, "Development of malaria diagnosis with convolutional neural network architectures: a CNN-based software for accurate cell image analysis". ITEGAM-JETIA, vol. 11, n°51, pp 35-42, 2025.
- [2] A. Jeneffa, E. Veemaraj, & A. Lincy, "ABM-OCD: Advancing ovarian cancer diagnosis with attention-based models and 3D CNNs", ITEGAM-JETIA, vol. 9, n° 43, pp. 23-33, 2023.
- [3] Y. Kumar, A. Koul, R. Singla, & M. F. Ijaz, M. F, "Artificial intelligence in disease diagnosis: a systematic literature review, synthesizing framework and future research agenda". Journal of ambient intelligence and humanized computing, vol. 14, n° 7, pp. 8459-8486, 2023.
- [4] M. S. Hashim, and A. A. Yassin, «Breast Cancer Prediction Using Soft Voting Classifier Based on Machine Learning Models», IAENG International Journal of Computer Science, vol. 50, n° 2, 2023.
- [5] F. AmrollahiI, S. P. Shashikumar, A. L. Holder et S. Nemati, «Leveraging clinical data across healthcare institutions for continual learning of predictive risk models» Scientific Reports, vol. 12, n° 11, p. 8380, 2022.
- [6] Q. An, S. Rahman, J. Zhou et J. J. Kang, «A comprehensive review on machine learning in healthcare industry: classification, restrictions, opportunities and challenges.» Sensors, vol. 23, n° 19, p. 4178, 2023.
- [7] M. Javaid, A. Haleem, R. P. Singh, R. Suman et S. Rab, «Significance of machine learning in healthcare: Features, pillars and applications,» International Journal of Intelligent Networks, vol. 3, pp. 58-73, 2022.
- [8] B. McMahan, E. Moore, D. Ramage, S. Hampson et B. A. y Arcas, «Communication-efficient learning of deep networks from decentralized data,» chez Artificial intelligence and statistics, PMLR, 2017, pp. 1273-1282.
- [9] H. A. P. K, M. R et A. S, «Jointly learning from decentralized (federated) and centralized data to mitigate distribution shift,» chez Proceedings of NeurIPS Workshop on Distribution Shifts, 2021.
- [10] M. De Lange, R. Aljundi, M. Masana, S. Parisot, X. Jia, A. Leonardis, G. Slabaugh et T. Tuytelaars, «A continual learning survey: Defying forgetting in classification tasks,» IEEE transactions on pattern analysis and machine intelligence, vol. 44, n° 17, pp. 3366-3385., 2021.

- [11] M. D. B. Thilagaraj, V. Pandimurugan, P. Naveen, M. S. Hema, S. Hariharasitaraman et P. Govindan, «A novel intelligent hybrid optimized analytics and streaming engine for medical big data,» *Computational and Mathematical Methods in Medicine*, vol. 2022, 2022.
- [12] X. Yang, H. Yu, X. Gao, H. Wang, J. Zhang et T. Li, «Federated Continual Learning via Knowledge Fusion: A Survey,» *IEEE Transactions on Knowledge and Data Engineering*, vol. 1, pp. 1-20, 2024.
- [13] M. A. Mohedano-Munoz, C. Soguero-Ruiz, I. Mora-Jiménez, M. Rubio-Sánchez, J. Álvarez-Rodríguez et A. Sanchez, «A streaming data visualization framework for supporting decision-making in the Intensive Care Unit,» *Expert Systems with Applications*, vol. 227, p. 120252, 2023.
- [14] Singh, M. B. Gurbuz, S. S. Gantha et P. Jasti, «Class-Incremental Continual Learning for General Purpose Healthcare Models,» *arXiv preprint arXiv*, vol. 2311, n° 104301, 2023.
- [15] M. Tilala, S. Pamulaparthivenkata, A. D. Chawda, and A. P. Benke, «Explore the Technologies and Architectures Enabling Real-Time Data Processing within Healthcare Data Lakes, and How They Facilitate Immediate Clinical Decision-Making and Patient Care Interventions,» *European Chemical Bulletin*, vol. 11, pp. 4537-4542, 2023.
- [16] Z. L. Teo, L. Jin, S. Li, D. Miao, X. Zhang, W. Y. Ng et D. S. W. Ting, «Federated machine learning in healthcare: A systematic review on clinical applications and technical architecture,» *Cell Reports Medicine*, 2024.
- [17] Jiménez-Sánchez, M. Tardy, M. A. G. Ballester, D. Mateus et D. Piella, «Memory-aware curriculum federated learning for breast cancer classification,» *Computer Methods and Programs in Biomedicine*, vol. 229, p. 107318, 2023.
- [18] M. Feng, Y. W. S. Yan, Y. Xu, L. Shao et H. Fu, «Specificity-preserving federated learning for MR image reconstruction,» *IEEE Transactions on Medical Imaging*, vol. 42, n° 17, pp. 2010-2021, 2022.
- [19] Liu, Cabezas, M. D. Wang, Z. Tang, L. Bai, G. Zhan, L. Y, K. Kyle, L. Ly, J. Yu, C. Shieh, A. Nguyen, K. E. Kandasamy, R. Sullivan, F. Calamante, M. Barnett, W. Ouyang, W. Cai et C. Wang, «Multiple sclerosis lesion segmentation: revisiting weighting mechanisms for federated learning,» *Frontiers in neuroscience*, vol. 17, n° 11167612, 2023.
- [20] H. Kassem, D. Alapatt, P. Mascagni, A. Karargyris et N. Padoy, «Federated Cycling (FedCy): Semi-Supervised Federated Learning of Surgical Phases,» *IEEE Transactions on Medical Imaging*, vol. 42, n° 17, pp. 1920-1931, 2023.
- [21] S. Pati, U. Baid, B. Edwards, M. Sheller, S. H. Wang, G. A. Reina et L. Poisson, «Federated learning enables big data for rare cancer boundary detection,» *Nature communications*, vol. 13, n° 11, p. 7346, 2022.
- [22] Qayyum, K. Ahmad, M. A. Ahsan, A. Al-Fuqaha et J. Qadir, «Collaborative federated learning for healthcare: Multi-modal covid-19 diagnosis at the edge,» *IEEE Open Journal of the Computer Society*, vol. 3, pp. 172-184, 2022.
- [23] L. Mondrejevski, Miliou, M. A. D. Pitts, J. Hollmén et P. Papapetrou, «Flicu: A federated learning workflow for intensive care unit mortality prediction,» *chez 2022 IEEE 35th International Symposium on Computer-Based Medical Systems (CBMS)*, Shenzhen, China, 2022.
- [24] T. K. Dang, X. Lan, J. Weng et M. Feng, «Federated learning for electronic health records,» *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 13, n° 15, pp. 1-17, 2022.
- [25] U. P. S. K. T. M. G. R. B. E. A. S. .. & B. Baid, «Federated learning for the classification of tumor infiltrating lymphocytes,» *arXiv preprint arXiv*, 2022.
- [26] L. Zhao et J. Huang, «A distribution information sharing federated learning approach for medical image data,» *Complex & Intelligent Systems*, vol. 9, n° 15, pp. 5625-5636, 2023.
- [27] M. Chetoui et M. Akhloofi, «Peer-to-Peer Federated Learning for COVID-19 Detection Using Transformers,» *Computers*, vol. 12, n° 15, p. 106, 2023.
- [28] W. Qiu, W. Ai, H. Chen, Q. Feng et G. Tang, «Decentralized federated learning for Industrial IoT with deep echo state networks,» *IEEE Transactions on Industrial Informatics*, vol. 19, n° 14, pp. 5849-5857, 2022.
- [29] T. V. Nguyen, M. A. Dakka, S. M. Diakiw, M. D. VerMilyea, M. Perugini, J. M. M. Hall et D. Perugini, «A novel decentralized federated learning approach to train on globally distributed, poor quality, and protected private medical data,» *Scientific Reports*, vol. 12, n° 11, p. 8888, 2022.
- [30] Z. Lian, Q. Yang, W. Wang, Q. Zeng, M. Alazab, H. Zhao et C. Su, «DEEP-FEL: Decentralized, efficient and privacy-enhanced federated edge learning for healthcare cyber physical systems,» *IEEE Transactions on Network Science and Engineering*, vol. 9, n° 15, pp. 3558-3569, 2022.
- [31] Li, H. Li et G. Yuan, «Continual Learning with Deep Neural Networks in Physiological Signal Data: A Survey,» *Healthcare*, vol. 12, n° 12, p. 155, 2024.
- [32] E. Batbaatar, K. H. Park, T. Amarbayasgalan, K. Davagdorj, L. Munkhdalai, V. H. Pham et K. H. Ryu, «Class-incremental learning with deep generative feature replay for DNA methylation-based cancer classification,» *IEEE Access*, vol. 8, pp. 210800-210815, 2020.
- [33] C. González, A. Ranem, D. Pinto dos Santos, A. Othman et A. & Mukhopadhyay, «Lifelong nnU-Net: a framework for standardized medical continual learning,» *Scientific Reports*, vol. 13, n° 11, p. 9381, 2023.
- [34] P. Kumari, J. Chauhan, A. Bozorgpour, R. Azad et D. Merhof, «Continual Learning in Medical Imaging Analysis: A Comprehensive Review of Recent Advancements and Future Prospects,» *arXiv preprint arXiv*, vol. 2312, n° 117004, 2023.
- [35] Y. Huang, C. Bert, S. Fischer, M. Schmidt, A. Dörfler, A. Maier et F. Putz, «Continual learning for peer-to-peer federated learning: A study on automated brain metastasis identification,» *arXiv preprint arXiv:2204.13591*, 2022.
- [36] Johnson, L. Bulgarelli, T. Pollard, L. A. Celi, R. Mark et S. Horng, «MIMIC-IV-ED (version 2.2),» *PhysioNet*, 2023.

- [37] F. Xie, J. Zhou, J. W. Lee, M. Tan, S. Li, L. S. O. Rajnther et N. Liu, «Benchmarking emergency department prediction models with machine learning and public electronic health records,» Scientific Data, vol. 9, n° 11, p. 658, 2022.
- [38] Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu et D. Warde-Farley, «Generative adversarial networks,» chez Advances in neural information processing systems, Montreal, Quebec, Canada, 2014.
- [39] L. Deng, X. He et J. Gao, «Deep stacking networks for information retrieval,» chez IEEE International Conference on Acoustics, Speech and Signal Processing, Vancouver, Canada, 2013.
- [40] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, & V. Smith, «Federated optimization in heterogeneous networks,» in Proceedings of Machine learning and systems, vol. 2, p. 429-450, 2020.
- [41] D. Qi, H. Zhao, and S. Li, «Better generative replay for continual federated learning,» arXiv preprint arXiv:2302.13001, 2023.