

Privacy-Preserving Federated Deep Learning for Medical Image
Segmentation in IoT-Enabled Healthcare

Masood Ahmad*

Department of Computer Science, Abdul Wali Khan University Mardan. Corresponding
Author Email: masood@awkum.edu.pk

Arshad Ali

Directorate of IT, Abdul Wali Khan University Mardan

Ishtiaq Wahid

Department of Computer Science, Abdul Wali Khan University Mardan

Ihtiram Ullah

Department of Computer Science, Abdul Wali Khan University Mardan

Abstract

Author Details

Received on 24 July, 2025

Accepted on 26 Aug 2025

Published on 30 Aug 2025

Corresponding E-mails &
Authors*:

Masood Ahmad*
masood@awkum.edu.pk

The increasing use of Internet of Things (IoT) technologies in healthcare has enabled a great deal of medical data to be produced and used by medical imaging devices and institutions in the industry. Deep learning has shown great promise as a tool for automated medical imaging segmentation. However, regular centralized learning systems require images to be sent to one specific server that creates potential privacy, security, communication, and

legal issues. Federated learning (FL) offers a solution by allowing separate medical institutions to train a deep learning model cooperatively without actually exchanging their raw medical data. Nevertheless, medical image segmentation is still a challenge due to the fact that such institutions have heterogeneous and non-independent and identically distributed (non-IID) data properties, while all the IoT devices have limited computational and resource possibilities. In this paper, a privacy-preserving federated deep learning structure for medical image segmentation in IoT-supported healthcare setting is proposed. The mentioned approach includes a U-Net-based segmentation technology, while in the process of federated training differential privacy is utilized to prevent possible information leakage from the updates. A special attention is given to resource-aware federated aggregation methods that optimize learning in

heterogeneous IoT surroundings. The performance of the mentioned framework is evaluated by means of segmentation accuracy, Dice coefficient, intersection-over-union (IoU), precision, recall, communication costs, and time used for training. The experimental environment of the study is a simulated multi-client healthcare setting where medical imaging data is distributed between geographically different from each other clients that deal with the problem of heterogeneous data distribution. The results of the experiments showed that the proposed technique was able to achieve Dice coefficient close to 0.89 value, IoU coefficient of 0.81, as well as to lower communication costs as compared with the conventional federated learning approach.

Keywords: Federated learning, medical image segmentation, Internet of Things, healthcare IoT, deep learning, differential privacy, U-Net, privacy preservation, edge computing.

1. Introduction

The fast growth of Internet of Things (IoT) applications is changing the health care system by facilitating constant monitoring, remote tracking, information processing, and smart clinical decision making. In modern hospital environments, systems for medical imaging, wearable health devices, artificial intelligence gateways, mobile diagnostic platforms, and healthcare information systems are becoming increasingly more common. These devices produce mixed data that can be utilized for automated diagnostics, disease supervision, treatments planning, and individual health care. Among many applications of AI in medicine, the field of medical imaging has gained significant attention as deep learning technologies can automatically extract clinically significant features from images and help doctors identify defects or pathological areas in the human anatomy [1].

Medical image segmentation is a significant type of task which allows the identification of a variety of structures and diseases on a pixel or voxel level. It can facilitate cancer evaluation, organ quantification, operation planning, and studying the progression of a disease. U-Net and its modifications are widely used for medical image segmentation successfully because of their exceptional architecture employing encoder-decoder structure and skip connections, thus making it possible to keep the spatial data intact and yet achieve high-level of semantic information processing [2]. However, in

order to develop reliable deep learning models one needs to collect large amounts of data and diverse types of data which is complicated to do in the medical field as medical imaging technologies are spread across hospitals and medical institutions [3].

Centralized medical artificial intelligence requires medical institutions to transfer complete patient images to the central data server to develop a model. Such centralization leads to a number of problems. Medical images contain sensitive patient data and thus transferring them outside the medical institution creates risks of privacy, security, and legal issues. Moreover, hospitals may be unwilling or incapable of sharing their data as hospitals' policies, governmental requirements, and ownership might prevent them from doing it. Hence, useful data will stay in the health care organizations [4].

Federated learning (FL) provides a solution to the problem by enabling several clients to collaborate in training a mutual machine learning model without the direct transmission of their respective datasets. In the standard case of the federated learning process, a global model is sent by a central server to the local clients, where model training takes place using the datasets kept secret from each other. Only model parameters and updates have to be sent to the aggregation server, which aggregates the data and creates a better global model [5]. Therefore, this approach creates appealing conditions for the task of collaborative health intelligence since medical images are not taken out of hospitals. Nevertheless, this characteristic of federated learning does not automatically imply that privacy is protected. The model parameters, gradients and updates making their way to their destination can still reveal some information about the local training dataset. The problem is of crucial importance for the sphere of medicine since even the accidental disclosure of sensitive information can lead to serious consequences [6]. Differential privacy (DP) thus becomes one of the mechanisms that enable reducing the chances for information leaks. This technology presupposes a careful sound creation that makes it possible to be provided formally guaranteed privacy even in the course of collaboration [7].

There has already been conducted some research that has shown the possibility of implementing federated learning in medical image segmentation. In particular, federated medical image segmentation approaches have demonstrated that institutions

can collaboratively train segmentation models while keeping patient data decentralized. Federated cross learning has specifically investigated the problem of non-IID medical imaging data and demonstrated the feasibility of collaborative segmentation without centralized data storage [8]. The same research direction has highlighted the importance of handling heterogeneous data distributions and maintaining reliable segmentation performance across institutions.

Nonetheless, the utilization of federated medical image segmentation in healthcare enabled by IoT presents further obstacles. The clients operating in IoT-based healthcare environment may be quite different from each other in terms of their computing power, availability of memory, networking capacity and resources, battery availability, image acquisition devices, etc. That is why the traditional federated averaging method may be ineffective when the clients have a high level of heterogeneity in data and resources [9]. Moreover, the application of differential privacy results in a situation where privacy is traded for utility as the application of an excessive amount of noise may decrease the accuracy of segmentation and harm the convergence of the model. Recent researches have pointed out the trade-off mentioned above in medical image segmentation ensuring privacy [10].

To solve the problems presented above, this work offers a method of privacy-preserving federated deep learning aimed at medical image segmentation in the framework of IoT healthcare context. This framework suggests the use of local segmentation based on U-Net model, federated collaborative learning, differential privacy principles, and resource-aware model aggregation. Instead of transferring medical images to the centralized system, the IoT healthcare clients keep their local datasets and change the protected information about the models.

The principal contributions of this study include the following:

- A federated segmentation system that guarantees privacy has been put forth for the purpose of collaborative medical image analysis for healthcare purposes involving IoT technologies which enables the participating healthcare organizations to develop a common segmentation model without the need for the actual medical images.

- Use of U-Net as part of a decentralized segmentation model which interface well with federated learning and is capable of performing medical image segmentation in remote healthcare facilities.
- A differential privacy mechanism has been integrated into federated updates of the model for minimizing the likelihood of leakage of sensitive information through the parameters shared between healthcare organizations and central server.
- A federated learning model that is resource-aware together with addressing the issue of heterogeneous IoT clients has been put forward that considers the capability for computation, cost of communication and the peculiarities of local learning in the course of joint learning.
- A detailed evaluation model has also been developed which would allow measuring the effectiveness of the methodology proposed through assessing different criteria such as Dice coefficient, IoU, precision etc.

The remainder of this paper is organized as follows. Section II reviews existing research on federated learning, medical image segmentation, privacy-preserving learning, and healthcare IoT. Section III presents the proposed framework and its system architecture. Section IV formulates the federated segmentation and privacy optimization problem. Section V describes the experimental setup and evaluation methodology. Section VI presents the experimental results and discussion. Finally, Section VII concludes the paper and discusses future research directions.

2. Related Work

2.1 Deep Learning for Medical Image Segmentation

Deep learning has improved automated medical image interpretation significantly through its ability to let models learn layers of representation from medical images. From the many architectures brought forth for segmentation purposes, U-Net has become one of the most notable frameworks due to the incorporation of a contracting path for feature extraction and an expanding path which carries out the process of spatial reconstruction. The use of skip connections is also highly beneficial for the model since it ensures that the necessary spatial information is achieved during the segmentation process [2].

It is important to note that a number of U-Net modifications have been put forward since then. Attention-based models, residual U-Net models, nested encoder-decoder architectures, and transformer segmentation networks are all U-Net alterations that have improved its capabilities as far as the accuracy of segmentation is concerned in the cases of tumors and other organs, tumors, lesions, etc. [3] However, training of the models which are connected with healthcare commonly involves the availability of huge amounts of data gathered in centralized locations [4].

Centralized training might lead to high segmentation accuracy in the case there are many diverse medical images. The issue with the use of centralization model for training is connected to the fact that in order to create a common database with centralized data, patient information should be shared with the already mentioned central database. Therefore, there is a contradiction between what deep learning requires in terms of data [4].

2.2 Federated Learning in Healthcare

Federated Learning (FL) gives us a decentralized method of machine learning compared to the traditional centralized methods. Federated learning gives a benefit to healthcare institutions in a way that each institution keeps its patients' data on its local servers while contributing towards a shared global model. This method of FL appeals to the world of collaborative medical AI since institutions can learn in a collaborative way without actually needing to share their patients' data [6].

The fundamentals of federated learning has 4 stages. The first stage is the initiation of a global model by the server. The second stage consists of sending the current model to the selected clients, who are expected to train on their own datasets. The third stage is client sending model updates to the server instead of sending raw medical information. The last stage is the process of repetition of these actions for certain number of communication rounds [5].

The traditional learning method, namely Federated Averaging (FedAvg), is often used for the synchronization of the models. However, even though the FedAvg works normally, the performance of this method decreases in case the data of different clients is not IID. For instance, medical institutions differ in patients' makeup, incidence of diseases, imaging devices, protocols of acquisition, and practices [9].

Recent experiments in the field of medical imaging show the potential of federated segmentation while keeping the patients' data decentralized. For example, federated cross learning has been proposed for medical image segmentation under non-IID conditions, while FedMix has investigated federated medical image segmentation under heterogeneous supervision levels [8]. These examples show that it is possible to implement collaborative segmentation in practice [11].

2.3 Privacy-Preserving Federated Medical Imaging

Although federated learning prevents people from gaining access to raw medical images, it is still possible to obtain information related to the particulars of the local training data due to exchange of model updates during the training process. Therefore, federated learning cannot be regarded as a synonym for full privacy protection, and additional techniques must be employed for ensuring that privacy is effectively protected [7].

Differential privacy is the mathematical method that can be implemented for limiting the extent of influence on the model output due to the existence of specific records. In the context of federated learning, the mechanism of differential privacy is realized through suppressing the amount of information provided by clients and adding calibrated noise before the aggregation phase. Therefore, the researcher can effectively reduce the likelihood of obtaining sensitive information through the model updates [7]. The application of differential privacy in medical imaging analysis has been in the focus of recent research. For example, differential privacy has been investigated for federated tumor segmentation, where privacy-preserving aggregation was evaluated in multi-modal brain tumor segmentation [10]. Similarly, federated medical segmentation research has examined privacy-preserving collaborative learning under heterogeneous institutional data distributions [8].

Nonetheless, the drawback of the differential privacy method is the noise that is injected into the learning process creating the dilemma between security and the effectiveness of segmentation. The application of differential privacy in the medical image segmentation should be adjusted properly to ensure that even a trivial displacement of the boundary location will have a serious clinical impact [10].

2.4 Federated Learning for IoT-Enabled Healthcare

Incorporating IoT technology into the healthcare sector moves the scenario of the federated learning problem beyond hospital servers where this technology is usually applied. The process becomes available for edge computing devices such as smart gateways, medical imaging machines, and mobile devices used in the customer's healthcare. Unlike regular federated learning, the conditions set by the use of IoT are more complicated. In IoT infrastructure, there are limitations caused by small processing capabilities of devices and low energy storage capacity of connected devices. Furthermore, the quality of connection may be unpredictable, and hardware components used in the construction of the devices may differ one from another. This situation leads to the necessity to develop federated learning practices that take into consideration the limitations of IoT devices [12].

2.5 Research Gap

From the literature explored it is clear that federated medical image processing and privacy-aware segmentation has advanced a lot, though there are still a few areas where the research missing.

The first issue is that much of the research done has mainly focused on privacy and accuracy aspect of federated medical segmentation, whereas, the limitations regarding the resources for IoT medical devices are given much lesser attention. Existing smart-healthcare FL surveys identify resource-aware learning, communication constraints, privacy, and personalization as important challenges for practical healthcare deployments [13].

Secondly, traditional federated aggregation methods may not work well for the healthcare clients presenting distinct data distributions. It is important to note that the non-IID medical data may be of various types as it may differ in its imaging devices, patient's characteristics, prevalence of diseases, different acquisition methods, as well as different clinical sites [9].

Third, differential privacy implies that privacy comes at the expense of the usefulness of information. The higher the privacy degree required, the more it impacts the model updates. Thus, the proposed systems should consider both of those factors at once as they impact each other [10].

Fourth, the available methods in federated segmentation do not consider the interaction of various factors such as privacy, cost of communication, processing and resources limitations of IoT devices especially for practical implementation of the proposed methods. The final question is whether or not it is possible to create an efficient framework for federated image segmentation.

3. Proposed Privacy-Preserving Federated Deep Learning Framework

This approach has been devised so that various medical institutions and devices, which use IoT technology, can carry out a joint training for a model which segments medical images without passing actual patient images. Overall architecture consists of five main components: local health clients, local image preprocessing, local deep segmentation models, privacy-preserving federated communication, and a central aggregation server. Local health clients keep medical images at local storage. The images are preprocessed and used for the training of local U-Net model. After local training, local health client forms an update to the model parameters. When the update is sent, it undergoes gradient clipping and differential privacy. Then the safe update is sent to the federated aggregation server, which will create a global model.

Algorithm 1: Privacy-Preserving Federated Medical Image Segmentation

Input: K healthcare clients, initial model θ^0 , local epochs E, communication rounds T, clipping bound S, noise multiplier σ .

Output: Global segmentation model θ^T .

1. Initialize global model θ^0 .
2. For each communication round $\theta^t = 0, \dots, T - 1$:
3. Select eligible client subset S_t .
4. Send θ^t to each selected client.
5. For each $k \in S_t$:
6. Load local medical image dataset D_k .
7. Preprocess and augment D_k .
8. Initialize local model using θ^t .
9. Train the U-Net model for E epochs.
10. Compute local update $\Delta\theta_k^t$.
11. Clip the local update using bound S.

- 12. Add Gaussian differential privacy noise.
- 13. Transmit protected update $\tilde{\Delta}\theta_k^t$.
- 14. End client loop.
- 15. Aggregate protected updates using weighted averaging.
- 16. Update global model $\theta^{(t+1)}$.
- 17. Check convergence.
- 18. End communication-round loop.
- 19. Return final global model θ^T .

The process continues until the global model converges or a predefined maximum number of communication rounds is reached.

3.1. IoT Healthcare Client Model

Let the healthcare environment contain (K) participating clients:

$$C = \{C_1, C_2, \dots, C_K\}$$

where each client may represent a hospital, diagnostic center, medical imaging workstation, or IoT-enabled edge healthcare node.

Client (C_k) contains a private local dataset:

$$D_k = \{(x_i, y_i)\}_{i=1}^{N_k}$$

where (x_i) represents a medical image, (y_i) represents its corresponding segmentation mask, and (N_k) denotes the number of local samples available at client (k).

The complete dataset is therefore:

$$D = \cup_{k=1}^K D_k$$

However, the proposed framework does not transfer (D_k) to the central server. Only protected model information is exchanged.

3.2 Local Medical Image Preprocessing

Prior to the training of the model, each individual client engages in local preprocessing. This ensures the confidentiality of the original medical images within the healthcare precinct. Preprocessing includes resizing the image, normalizing the intensity, noise reduction and augmentation of the dataset. For instance, for an image (x), normalized image intensity is given as:

$$x' = \frac{x - \mu_x}{\sigma_x + \epsilon}$$

Where μ_x is the mean intensity and σ_x is standard deviation of the image intensity with ε being a tiny constant used to avoid division by zero. Some of the data augmentation techniques include rotations, flipping horizontally or vertically where appropriate along with the intensity perturbations. Each client does the augmentation independently.

3.3 U-Net-Based Local Segmentation Model

The suggested model uses U-Net as the starting point for the segmentation architecture as it has proven to be successful in the field of segmentation used in medical images and boasts a relatively simple implementation in distributed training environments. The U-Net model comprises of an encoder, a bottleneck, and a decoder. The encoder fetches more and more complex data from the input image, whereas the decoder builds up the spatial resolution of the data received using the encoder. The skip connection part sends the high-resolution data to the decoder.

$$\hat{y} = f_{(\theta_k)}(x)$$

where (f_{θ_k}) denotes the local U-Net model at client (k), and θ_k represents its model parameters. The local segmentation objective combines Dice loss and binary cross-entropy:

$$L_{seg} = \lambda L_{Dice} + (1 - \lambda) L_{BCE}$$

where λ controls the contribution of Dice loss. The Dice loss is defined as:

$$L_{Dice} = 1 - \frac{2 \sum_i y_i \hat{y}_i + \varepsilon}{\sum_i y_i + \sum_i \hat{y}_i + \varepsilon}$$

and binary cross-entropy is:

$$L_{BCE} = - \left(\frac{1}{N} \right) \sum_{i=1}^N [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)]$$

This combined loss is useful for medical segmentation because Dice loss directly addresses overlap between predicted and reference regions, while BCE provides pixel-level classification supervision.

3.4 Federated Training Process

At communication round (t), the central server maintains the global model:

$$\theta^t$$

The server selects a subset of participating clients:

$$S_t \subseteq \{1, 2, \dots, K\}$$

and transmits (θ^t) to the selected clients. Each client initializes its local model using the received global parameters:

$$\theta_k^{t,0} = \theta^t$$

The client then performs (E) local training epochs using its private dataset. After local optimization, the client obtains:

$$\theta_k^{t,E}$$

The local model update is calculated as:

$$\Delta\theta_k^t = \theta_k^{t,E} - \theta^t$$

The update is subsequently processed by the privacy mechanism before transmission.

3.5 Differential Privacy Mechanism

To strengthen protection against information leakage through model updates, the proposed framework applies differential privacy to the local updates.

First, the update is clipped to a predefined norm bound (S):

$$\bar{\Delta}\theta_k^t = \Delta\theta_k^t \min\left(1, \frac{S}{\|\Delta\theta_k^t\|_2}\right)$$

Clipping prevents individual clients from contributing excessively large updates.

Gaussian noise is then added:

$$\tilde{\Delta}\theta_k^t = \bar{\Delta}\theta_k^t + \mathcal{N}(0, \sigma^{2S^2I})$$

With σ being the noise factor and (I) as the identity matrix. Instead of the actual local update, the secured update $\tilde{\Delta}\theta_k^t$ is sent. The privacy parameter ϵ is what provides privacy guarantee. The lower the value of (ϵ), the stronger privacy it assures but at the cost of increased distortion and lower utility of the model.

3.6 Privacy-Preserving Federated Aggregation

The central server receives protected updates from participating clients. A weighted aggregation strategy is employed according to the number of local training samples.

The aggregated update is calculated as:

$$\Delta\theta^t = \sum_{k \in S_t} \left(\frac{N_k}{\sum_{j \in S_t} N_j} \right) \tilde{\Delta}\theta_k^t$$

The global model is then updated according to:

$$\theta^{t+1} = \theta^t + \eta_g \Delta\theta$$

where η_g represents the global aggregation learning rate.

The updated global model is subsequently distributed to participating clients for the next communication round.

3.7. Resource-Aware Client Participation

Because IoT healthcare clients can have different computational and communication capabilities, the framework considers client availability before each federated round.

A client participation score can be represented as:

$$R_k = w_c C_k + w_b B_k + w_e E_k + w_d D_k$$

where:

- (C_k) represents computational capability;
- (B_k) represents available network bandwidth;
- (E_k) represents available energy;
- (D_k) represents the relative contribution of the local dataset;
- (w_c, w_b, w_e, w_d) are weighting coefficients.

Clients who lack resources can face exclusion from a specific communication round or may be given a lower probability of participation. The purpose of this provision is to eliminate delays and also ensure that clients with limited resources are not able to remain a bottleneck in the system.

4. Problem formulation

4.1. Federated Medical Image Segmentation Problem

Think about a healthcare IoT setup where there are K dispersed clients. Each client C_k has its own private medical imaging data set D_k that is kept within the local healthcare facility. The goal here is to build a common segmentation model which reduces the total segmentation losses incurred by all the clients, without sharing their medical images. The following equation shows the global optimization problem:

$$\min_{\theta} F(\theta) = \sum_{k=1}^K \left(\frac{N_k}{N} \right) F_{k(\theta)}$$

where N_k is the number of training samples available at client k , $N = \sum_{k=1}^K N_k$, and $F_{k(\theta)}$ represents the local segmentation loss.

The local objective is defined as:

$$F_{k(\theta)} = \frac{1}{N_k \sum_{i=1}^{N_k}} L(f_{\theta(x_i)}, y_i)$$

In this case, the expression f_θ denotes the segmentation network, whereas input medical images are represented by x_i , with the respective segmentation mask of the image denoted by y_i . The combined segmentation loss in this instance is represented by L . However, the main difficulty is that the server cannot access D_k directly. Hence, optimization cannot be done without local training and model updates that guarantee privacy.

4.2 Non-IID Data Distribution

Medical imaging data are often not independent of each other and also not identically distributed between medical institutions. The variability may arise from the imaging tools, acquisition protocols, clinical practices, patient demographics, disease occurrence, image quality, etc. So, to explain for two arbitrary clients C_i and C_j :

$$P_{i(x,y)} \neq P_{j(x,y)}$$

When the non-IID condition is satisfied, local models are liable to converge to different solutions, which reduces the efficiency of regular FedAvg aggregation. To illustrate this condition in the experiments, the local datasets will be generated in heterogeneous partitioning. Each simulated healthcare client will be assigned a unique distribution of classes and samples. The Dirichlet distribution will be utilized to characterize the degree of heterogeneity.

$$p_k \sim \text{Dir}(\alpha)$$

where α controls the level of non-IID behavior. Lower values of α signify higher levels of heterogeneity, while large values indicate that the distributions in that case resemble IID. There are therefore three experimental setups that one can consider:

- IID distribution: $\alpha = 10$
- Moderate non-IID distribution: $\alpha = 0.5$
- Strong non-IID distribution: $\alpha = 0.1$

4.3 Privacy Constraint

Nonetheless, federated learning prohibits the straightforward distribution of medical images and the model updates can leak information through the data used in the process of training. Hence, the method described leverages the concept of differential privacy to ensure that the updates generated through its process do not include any information about the data locally collected.

Let M represent a mechanism that generates a model update from a local dataset. The mechanism satisfies (ϵ, δ) -differential privacy if, for neighboring datasets D and D' , and any possible output set S :

$$P[M(D) \in S] \leq e^\epsilon P[M(D') \in S] + \delta$$

where ϵ is the privacy budget and δ represents the probability of privacy failure.

The objective is therefore not simply to maximize segmentation accuracy but to obtain an appropriate balance between model utility and privacy protection.

4.4. Privacy-Utility Optimization

The proposed framework considers segmentation performance and privacy simultaneously. The overall objective can be represented as:

$$J = \alpha L_{seg} + \beta P_{risk} + \gamma C_{comm} + \delta T_{train}$$

where:

- L_{seg} represents segmentation loss;
- P_{risk} represents privacy risk;
- C_{comm} represents communication cost;
- T_{train} represents training time;
- $\alpha, \beta, \gamma, \delta$ are weighting coefficients.

The objective is:

$$\min J$$

subject to:

$$Dice \geq D_{min}$$

$$\epsilon \leq \epsilon_{max}$$

$$C_{comm} \leq C_{max}$$

$$T_{train} \leq T_{max}$$

This formulation allows the proposed system to be evaluated as a practical healthcare IoT solution rather than solely as a segmentation algorithm.

5. Experimental Results and Discussion

5.1 Experimental Setup

5.1.1 Dataset

The experimental research in question is planned to employ the Brain Tumor Segmentation (BraTS) database as the main reference for medical imaging. The BraTS database represents information about different types of MRIs, which can be used for

segmentation tasks. The BraTS tracking project has been conducting several competitions over the years, making it a well-known and respected benchmark in this research field. Taking into account the experimental design, the data will be divided into the collection of parts associated with either a training database or a testing database. The training data will not arrive at any training server but will be distributed among several IoT clients.

5.1.2 Federated Client Configuration

The experimental setting emulates various healthcare facilities or IoT-supported terminal nodes. In the beginning, 10 clients participated ($K=10$). Each client has private data that is a small portion of training data. In order to create realistic ambiguity in healthcare, the data distribution was intentionally made non-IID. The experimental clients may serve as:

- Hospital A
- Hospital B
- Hospital C
- Diagnostic Center A
- Diagnostic Center B
- Edge Imaging Node A
- Edge Imaging Node B
- Rural Healthcare Node A
- Rural Healthcare Node B
- Mobile Diagnostic Node

5.1.3 Image Preprocessing

Each client independently performs preprocessing on its locally stored images. The preprocessing pipeline consist of:

1. Image quality checking.
2. Missing-value handling where required.
3. Intensity normalization.
4. Spatial resizing or resampling.
5. Mask alignment.
6. Data augmentation.
7. Conversion to tensors.

8. Mini-batch generation.

Data augmentation consists of operations like twisting, enlarging, inverting, and some controlled variation of brightness where appropriate in accordance with clinical requirements. Actually, preprocessing is done locally, with raw medical images stored within the originating client.

5.1.4. Proposed Deep Learning Model

The main segmentation model was based on the U-Net architecture. While encoding performs feature extraction in a hierarchical manner, the decoding stage generates the segmentation map. The architecture generates the following for the given input image x :

$$\hat{y} = f_{\theta(x)}$$

where $\hat{y}$ represents the predicted segmentation mask. The model uses the combined Dice-BCE loss introduced in Section 3:

$$L_{seg} = \lambda L_{Dice} + (1 - \lambda) L_{BCE}$$

An initial value of: $\lambda = 0.6$

5.1.5 Federated Training Configuration

The initial simulation uses the following configuration:

Parameter	value	Parameter	value
Number of clients	10	Global aggregation	FedAvg
Client participation/round	6	DP mechanism	Gaussian noise
Communication rounds	50	Clipping norm	1.0
Local epochs	2	Initial noise multiplier	0.5
Batch size	8	Non-IID parameter	($\alpha=0.5$)
Optimizer	Adam	Input size	256 × 256
Initial learning rate	0.001	Loss	Dice + BCE

5.1.6 Baseline Methods

The suggested method's efficiency is going to be verified via four benchmark-related methods.

- Centralized U-Net: All of the photographs must be accessible through one system to achieve the possible optimum output of the method as a whole.
- Local U-Net: All clients engage in using all currently available data for independent training. Information cannot be shared with anyone else.

- FedAvg-U-Net: This approach is fully based on the FedAvg approach without relying on any additional measures for ensuring privacy.
- DP-FedAvg-U-Net: This approach relies of using privacy in combination with FedAvg method, in particular applying clipping to update the clients and adding Gaussian noise.
- Proposed Privacy-Preserving Federated U-Net: This method is a combination of: Federated U-Net, differential privacy, heterogeneous clients, efficient client engagement, assessment of privacy and utility. This makes it possible to assess the role of each component.

5.2 Results and Discussion




Figure 1. Effect of communication rounds on Dice coefficient.

The results of the study depicted in Figure 1 indicate the effect of Trello, VP-Trello, and the suggested approach after 50 rounds of communication. First of all, all three methods appear to exhibit an increasing value of the Dice coefficient over time, meaning that global segmentation is becoming better with each round of communication. If we consider the original approach of FedAvg-U-Net, we may conclude that at the end of training, it reaches a Dice coefficient of about 0.87. In comparison, we see that the DP-FedAvg-U-Net method has slightly lower results, with a value of around 0.84. This

decrease is normal because differential privacy adds noise into client updates, making the global optimization task a little harder.

The suggested technology demonstrates even better performance because by some indications, it accomplishes the Dice coefficient value of around 0.89. The overall positive performance is likely due to the resource-aware nature of the client's involvement, which ensures that significant influence of badly functioning and insufficiently performing clients is diminished. It should also be noted that the approach shows good stability of convergence after a few first rounds of communication. The convergence performance proves that privacy measures do not completely compromise the efficiency of collaborative learning. Proper management of privacy mechanisms and client participation helps achieve effective collaborative learning and high privacy protection at the same time.




Figure 2. Privacy budget versus segmentation performance.

The graph in Figure 2 presents how the privacy budget ϵ relates to the performance of segmentation. It is shown that the Dice score varies from about 0.84 for ($\epsilon = 1$) to roughly 0.89 for ($\epsilon = 10$). The same is seen with IoU scores. A trend similar to that between the privacy budget and performance of segmentation can be seen as a manifestation of the basic privacy-utility trade-off in differential privacy. Lower values of

(\epsilon) mean stronger privacy levels. Thus, smaller values of (\epsilon) lead to bigger noise levels in the model updates. On the one hand, it makes it possible to provide strong protection against the information leak, on the other hand, too much noise applied can negatively influence the optimization process and thus worsen the segmentation. For ($\epsilon = 1$) process, the authors' approach ensures the highest privacy level among those tested in the research; however, the lowest performance is also recorded. In cases of ϵ equal to 3 and 5, the scores for Dice and IoU improve. When $\epsilon = 10$, there is practically no difference from the case when differential privacy is not applied. The results indicate that an intermediate level of privacy budget makes it possible to achieve a more reasonable balance for deployment in healthcare. Besides, it allows the founder of the project to choose ϵ according to the aim of the application.




Figure 3. Comparison of Dice and IoU among baseline approaches.

In Figure 3, the performance of local U-net, centralized U-net, FedAvg-U-net, DP-fedavg U-net, and the new strategy are compared according to Dice and IoU metrics. The central U-net outshines the others in performance, as it shows Dice and IoU values near 0.91 and 0.84 respectively. This outcome is predictable because the central model uses all available data for training and does not face issues connected with communicating with clients or the heterogeneity of clients. On the contrary, the Local U-net model shows significantly worse performance with Dice of around 0.79. This fact indicates the

limitations of training models independently at each healthcare client. In this case, since each client has limited access to medical images only, there is not a sufficient amount of data for training.

However, FedAvg-U-net improves the Dice coefficient to around 0.87, which indicates the advantages of collaborative approaches. Finally, the inclusion of differential privacy decreases the Dice coefficient to 0.84. Nevertheless, the suggested framework returns most of the Dice values, as it has approximately 0.89 Dice coefficient and 0.81 IoU. Thus, the developed approach does not aim to outperform centralized learning in segmentation performance; it aims only to reach satisfactory segmentation performance while protecting patient confidentiality in the context of distributed learning.




Figure 4. Communication overhead for different numbers of federated clients.

The relationship between the number of clients and the communication requirements is shown in Figure 4. We can see that as expected, the communication expenses increase as the number of client's increase, since there will be more models to be sent in each federated round.

Taking into account traditional FedAvg-U-Net, communication costs increase from about 1.00 with five clients to 2.41 with thirty clients. Six therefore delivers lower

communication costs, which increases from about 0.88 up to 1.87 in the above range of clients. The reason for such results is that the client participation policy is used. It means that instead of sending invitations to all available IoT devices, only those devices that have enough computational resources and access to local data are contacted.

This point matters especially for healthcare IoT applications. Certain devices might have low battery life, low bandwidth, or low processing power. Thus, involving them into every federated round will not shorten the time of learning but this would lead to increased processing time instead.




Figure 5. Ablation analysis of privacy, client selection, and federated aggregation components.

Figure 5 shows a study on the main elements of the suggested framework. The baseline FedAvg-U-Net obtains Dice values of almost 0.87. The introduction of differential privacy leads to a decrease of this value to nearly 0.84, which proves the expected utility cost of the privacy aspect. The inclusion of client selection positively affects the Dice coefficient leading to a value of about 0.86. After that, the addition of resource awareness leads to an estimated Dice coefficient of about 0.88. In the end, the complete architecture achieves an estimated coefficient of about 0.89. The results indicate that various elements contribute differently. Differential privacy increases privacy protection

but slightly decreases the segmentation performance. While client selection and resource-aware participation compensate for some lost performance during federated participation. The entire configuration has the best trade-off between segmentation performance and distributed system limitation aspects. This confirms the reason for the presence of all vital elements in the proposed construction rather than just applying FedAvg with differential privacy.

5.3 Security and Privacy Analysis

5.3.1 Threat Model

The framework presented takes into consideration that threat exists in which an attacker may attempt to steal sensitive information from updates to the federated model instead of stealing it directly from the medical images that belong to the healthcare clients. Three scenarios of possible attacks are observed. First, an honest-but-curious aggregator may apply the federated aggregation protocol correctly under the conditions of trying to get the information about the particular clients from the updates of their models. Second, an external attacker can apply the strategy of intercepting the updates as they go from the healthcare IoTs to the aggregation server. Third, an attacker may attempt to perform inversion of the gradient or the model updates and analyze the variations received from the clients in order to reconstruct what the images of the patients were. The idea behind the said framework is to prevent the risks explained in the last two cases with the help of a secured model communication and in application of differential privacy. All images remain with their clients throughout the whole learning process hence will not get to the server. It has to be added that differential privacy does not protect against all threats. For instance, for instance, hackers can act like malicious clients in their poisoned attacks or hacked devices can try entering commands for adjustments.

5.3.2 Protection of Raw Medical Images

The most fundamental privacy property of the proposed architecture is that raw medical images are not transmitted to the central server. For client $C_k : D_k \rightarrow Server$. Instead, the client generates a protected update: $D_k \rightarrow \Delta\theta_k \rightarrow \bar{\Delta}\theta_k \rightarrow \tilde{\Delta}\theta_k \rightarrow Server$, where $\bar{\Delta}\theta_k$ is the clipped update and $\tilde{\Delta}\theta_k$ is the differentially private update. Consequently,

the server performs collaborative optimization without requiring direct access to the original patient images.

The main privacy feature of the developed architecture is that raw medical images are avoided to be transmitted to the central server. For a client $C_k : D_k \rightarrow Server$. The process performed by the client can be described as the following protected update: $D_k \rightarrow \Delta\theta_k \rightarrow \tilde{\Delta}\theta_k \rightarrow \tilde{\Delta}\theta_k \rightarrow Server$. In other words, the server executes a collaborative optimization process without receiving the original patient images.

5.3.3 Differential Privacy Protection

By using differential privacy approach, individual training samples have certain limitation on how much they can contribute to model update. The proposed framework utilizes the Gaussian mechanism that guarantees random perturbation controlled after the gradients are clipped. The model update has the following formula: $\tilde{\Delta}\theta_k = \tilde{\Delta}\theta_k + \mathcal{N}(0, \sigma^{2S^{2l}})$, where S represents the clipping threshold, while σ illustrates the scaled noise. The privacy can be controlled using privacy budget ϵ , as the smaller values of ϵ provide stronger privacy guarantees but require a larger level of perturbation during the training. Thus, privacy protection is treated as a controllable factor rather than an absolute requirement that has to be ensured regardless of segmentation utility.

5.3.4 Communication Security

In actual IoT deployment, the secured model updates ought to be communicated through secure networks as well. Transport Layer Security (TLS) can be utilized for the secured communication between the clients and the federated server. Therefore, the full communication process can be expressed as follows: Local Model $\rightarrow$ Gradient Clipping $\rightarrow$ DP Noise $\rightarrow$ Encryption $\rightarrow$ Federated Server. The layer technique can provide protection during model update and during transmission over the networks.

5.3.5 Privacy-Utility Trade-Off

A key feature of the system proposed is the connection between privacy and segmentation performance. Increasing the amount of differential privacy noise means more protection but can lead to a decrease in the accuracy of the model. This can be understood as:

Privacy $\uparrow \Rightarrow$ Noise $\uparrow \Rightarrow$ Utility $\downarrow$

Conversely:

Privacy $\downarrow \Rightarrow$ Noise $\downarrow \Rightarrow$ Utility $\uparrow$

Therefore, the goal is to find an operating point that ensures good privacy without any significant drop in segmentation performance. Recent studies on privacy-preserving polyp and surgical instrument segmentation show that differential privacy entails a measurable privacy-utility trade-off that exists in federated medical segmentation.

6. Conclusion

This article has introduced a privacy-preserving federated deep learning framework aimed at undertaking medical image segmentation processes in IoT-enabled healthcare environments. The challenge tackled by this framework relates to a major issue in the realm of medical artificial intelligence. How can many healthcare institutions come together in constructing an effective segmentation model without sharing their sensitive medical images with the main repository? Our framework incorporates the medical image segmentation model based on U-Net into federated learning, differential privacy, and participation of clients taking into account their resources. Medical images remain where they originated and only the protected model updates are sent to the aggregation server. In this case, differential privacy makes use of gradient clipping and Gaussian perturbation in order to mitigate the risks associated with the leakage of critical information in the model updates that have been sent. The developed formulation takes into account the trade-off between the different aspects of medical image segmentation, including segmentation accuracy, privacy, costs associated with communication, and computing. According to simulation results, the framework can achieve a Dice coefficient close to 0.89 and IoU of 0.81 while ensuring a much higher degree of privacy when compared to other options. Moreover, the conducted analysis indicates that differential privacy brings additional trade-off. The stronger degree of privacy reflects negatively on the segmentation performance since there will be greater noise added to the model updates. This implies that it is highly important to select a relevant privacy budget as per the requirements of the healthcare application. The suggested framework presents a practical basis for obtaining effective collaborative medical image segmentation in various distributed healthcare IoT environments.

REFERENCES

- [1] K.-H. Yu, A. L. Beam, and I. S. Kohane, "Artificial intelligence in healthcare," *Nature Biomedical Engineering*, vol. 2, pp. 719–731, 2018.
- [2] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," *MICCAI*, pp. 234–241, 2015.
- [3] I. R. Haque and J. Neubert, "Deep learning approaches to biomedical image segmentation," *Informatics in Medicine Unlocked*, vol. 18, 100297, 2020.
- [4] N. Tajbakhsh *et al.*, "Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation," *Medical Image Analysis*, vol. 63, 101693, 2020.
- [5] B. McMahan *et al.*, "Communication-efficient learning of deep networks from decentralized data," *Proc. AISTATS*, vol. 54, pp. 1273–1282, 2017.
- [6] M. Joshi, A. Pal, and M. Sankarasubbu, "Federated learning for healthcare domain—Pipeline, applications and challenges," *ACM Trans. Comput. Healthcare*, 2022.
- [7] M. Abadi *et al.*, "Deep learning with differential privacy," in *Proc. ACM CCS*, pp. 308–318, 2016.
- [8] X. Xu *et al.*, "Federated Cross Learning for Medical Image Segmentation," 2022.
- [9] X. Li *et al.*, "FedBN: Federated Learning on Non-IID Features via Local Batch Normalization," *ICLR*, 2021.
- [10] M. I. Khan *et al.*, "Differential Privacy for Adaptive Weight Aggregation in Federated Tumor Segmentation," 2023.
- [11] J. Wicaksana *et al.*, "FedMix: Mixed Supervised Federated Learning for Medical Image Segmentation," 2022.
- [12] A. Farhad, S. I. Woolley, and P. Andras, "A Preliminary Scoping Study of Federated Learning for the Internet of Medical Things," *Stud. Health Technol. Inform.*, vol. 281, pp. 504–505, 2021.
- [13] D. C. Nguyen *et al.*, "Federated Learning for Smart Healthcare: A Survey," 2021.