Journal of Neonatal Surgery
ISSN(Online): 2226-0439
Vol. 14, Issue 25s (2025)
https://www.jneonatalsurg.com

OPEN ACCESS

Federated Learning for Brain Tumor Segmentation in MRI Scans: A Privacy-Preserving and Domain-Adaptable Approach

Partheepan R[1], S.S.Sivaraju[2], A.Senthilkumar[3], V.Radhika[4] , N.Suba Rani[5], S.Sivarajan[6]

[1]Associate Professor Department of ECE Velammal Institute of Technology.
Email ID: rpartheepanphd@gmail.com
[2]Professor Department of EEE RVS College of Engineering & Technology.
Email ID: sssivaraju@gmail.com
[3]Professor Department of CSE DayanadaSagar University, Bangalore.
Email ID: senthil.kumar-cse@dsu.edu.in
[4]Associate Professor, Department of Biomedical Engineering Sri Ramakrishna Engineering College.
Email ID: radhika.senthil@srec.ac.in
[5]Associate Professor, Dr.Mahalingam College of Engineering and Technology.
Email ID: suba.journal@gmail.com
[6]Assistant Professor Department of EEE Vel Tech Multi Tech Dr.Rangarajan Dr.sakunthala Engineering college.
Email ID: ssivarajan78@gmail.com

Cite this paper as: Partheepan R, S.S.Sivaraju, A.Senthilkumar, V.Radhika, N.Suba Rani, S.Sivarajan, (2025) Federated Learning for Brain Tumor Segmentation in MRI Scans: A Privacy-Preserving and Domain-Adaptable Approachs. *Journal of Neonatal Surgery*, 14 (25s), 766-776.

ABSTRACT

The application of deep learning in intricate tasks like brain tumour segmentation heavily relies on potent algorithms, which in turn depend on large and heterogeneous datasets that are often fraught with privacy issues or deep domain variations (domain shift). Contra-violent methods concerning MRI data hold little promise as sharing sensitive data captured from patients across various institutions breaches privacy laws like HIPAA; furthermore, variation in scanner equipment and its associated detailing parameters lead to a pathological model building which does not generalise well on new data.

This document posits a robust approach to tackle these issues using Federated Learning (FL). With FL, it's possible to train a model in collaboration with various institutions without sharing raw data from patients. In contrast, model training takes place privately first at each institution, and only learned model updates (i.e., the parameters or gradients of the model) are sent for aggregation on a central server. This allows for the refinement of a global model while keeping patient data safe within their institution. In addition, FL provides and preserves anywhere the variance of medical imaging data within one defined place as compared to traditional approaches, boosting the model resilience and flexibility to cope with new challenges due to changing MRI protocol shifts.

This approach included first locally training a Convolutional Neural Network (CNN) on data partitioned at each client site and then aggregating the model weights at a central location. The model was evaluated at the client based on the loss incurred during training and accuracy during prediction, achieving the intended goals.

Despite experiencing certain issues such as data heterogeneity (as shown by one client's loss oscillation), the model was able to learn from the distributed data and predict tumours correctly on the MRI scans. This study concludes that federated learning provides an effective framework with privacy-preserving characteristics and flexible domain adaptation for constructing resilient brain tumour segmentation models with multiple decentralised MRI datasets.

1. INTRODUCTION

One of the most important works of neuro- oncology is the correct extraction of a brain tumour from an MRI scan. Correct segmentation is not an amateur effort; it also defines, at a very fundamental level, the critical clinical decisions that define accuracy in diagnosis, during treatment plan formulation, and while monitoring disease status whether it is evolving or regressing. Recently, the domain of medical image analysis has undergone a sharp change owing to the introduction of deep learning (DL). In particular, convolution neura network (CNN) architectures, including the U-Net and its many permutations, have successfully automated the intricate task of tumour segmentation and are now recognised to perform comparably to human experts if not better.

Even with these advancements, powerful DL models have yet to be adapted into robust clinical tools readily available in hospitals and clinics. The capability for, and almost all of the ability to generalise meaning perform accurately on new data not seen during training are almost completely reliant on having access to multiple heterogeneous datasets that are exhaustively annotated. Building any kind of dataset, especially for critical healthcare needs such as brain tumour analysis, is difficult.

There are extreme constraints on the handling of patient data, such as the HIPAA regulations in the United States, enforcing the utmost secrecy concerning medical images and their metadata between institutional boundaries.

The first issue limits the creation of comprehensive models as the repository systems necessary are centrally located and highly regulated. The second issue regarding image annotation is possibly the most complex as it needs to be completed by specialised neuro-oncologists or trained radiologists. This implies the process will be expensive and significantly time-consuming. Hence the datasets available are scant.

Third but still equally significant is domain shift which is caused by the variability in MRI acquisition protocols. Imaging centres and hospitals use scanners from a variety of manufacturers, each with unique magnetic field strengths, as well as differing parameters such as resolution, contrast, FLAIR, T1, and T2. Such variation causes characteristics such as image quality, contrast, and texture within images sourced from different datasets to differ greatly. For instance, a model trained using protocols from a specific hospital tends to perform sub-optimally with data from other domains making the model less functional in real life scenarios.

To tackle the intersection of scarce data, privacy concerns, and heterogeneous data, this paper leverages Federated Learning (FL). FL is a representation of a new era in distributed machine learning as it allows multiple stakeholders like hospitals to collaboratively train a posteriori shared prediction model without locally storing the pooled data or sensitive raw-information.

Within the FL framework, model training is performed at the local level on data collected from each participating institution. Instead of sending private MRI scans, only model updates—parameter weights or gradients, for example—are exchanged periodically to a coordinating server. This server aggregates participant updates, frequently employing a weighted averaging technique, to improve the global model, which is subsequently sent back to the institutions for another round of local training. This recurring cycle permits the global model to ingest information from the numerous distributed datasets while maintaining the patient's data securely within the institution's boundaries, thus preserving privacy without additional efforts. Furthermore, FL is ideally designed to manage the non-IID (non-identically and independently distributed)data characteristics typical in multi-institutional scenarios, presenting an opportunity to develop models that better withstand and adapt to the domain shifts caused by varying MRI protocols. Hence, this study proposes and explores the applicability of FL towards constructing efficient deep learning models for brain tumour recognition that are privacy friendly and cross-domain adaptable to varying MRI data.

2. KEYWORDS

Federated Learning, Brain Tumor Segmentation, Magnetic Resonance Imaging (MRI), Deep Learning, Convolutional Neural Network (CNN), Data Privacy, Domain Adaptation, Medical Imaging.

3. RELATED WORK

Within the scope of clinical practice, the segmentation of brain tumours through Magnetic Resonance Imaging (MRI) poses a significant challenge considering issues like data diversity and confidentiality. This part of the paper summarises the literature on brain tumour segmentation with a particular emphasis on the development of deep learning techniques, domain adaptation, data scarcity solutions, and the rise of federated learning.

MRI-based Brain Tumour Segmentation

Deep learning techniques, particularly with Convolutional Neural Networks (CNNs), fuel automated brain tumour segmentation. Ullah et al. suggested a hybrid technique that employs CNNs and also enhances segmentation through intensity, texture, and shape-based handcrafted feature extraction [8]. In the MRI modality undersampling problem, Zhou et al. proposed improving segmentation using non-available scans with latent correlation modelling in Latent Correlation Representation Learning (LCRL) [17]. Li's research also aimed at tackling the incomplete multi-modal MRI scans problem using a knowledge transfer network [7]. Menon and Ramakrishnan went beyond deep learning and researched unsupervised segmentation using Artificial Bee Colony (ABC) algorithms and FCM clustering [13]. Another angle is enhancing the MRI itself; Zhou et al. presented MRBT-SR-GAN, which super-resolves MR images of brain tumours [19]. Other contributions come from traditional image processing, such as the work of Nakhmani et al. who employed adaptive Sobolev snakes for semi-automatic segmentation and necrosis detection [16]. The importance of accurate segmentation for predicting patient survival outcomes and life quality is well demonstrated by Sun et al.

Based on the citation, who applied deep learning techniques for segmentation and survival prognosis on multimodal scans is given in [12]. Systematic reviews look at the approaches or directions of AI (ML and DL) integration to MRI brain tumour diagnostics over a long time such as Satushe et al.'s work [3]. Moreover, Verma and Yadav focused on reviewing neuroimaging and deep learning networks dedicated to brain tumour segmentation [18]. Also, Munir et al. designed and analysed hybrid deep learning models using CNNs and inception modules which they proposed and compared [20]. Trust in automated systems led ND et al. to suggest incorporating XAI techniques such as SHAP and LIME into a deep learning model for explainable segmentation [14].

Domain Shift and Generalisation in Medical Imaging

The adaptation to clinical use cases is often limited by the variability of deep learning models to changes in the scanning device, protocols used for data collection, or the sample population. This change usually requires more attention on endurance and general clarity [10]. To mitigate loss of performance with DA, an attempt is made to close the difference of starting and end goals or domains. Zhang et al.

We proposed MAPSeg, a framework for Unsupervised Domain Adaptation (UDA) using masked autoencoding and pseudo-labelling for heterogeneous medical image segmentation across settings such as centralised, federated, and test-time UDA [5]. Domain Generalisation (DG) aims at an even more challenging problem of generalising to wholly unseen domains. Yanzhen et al. reviewed the literature focusing on cross-domain difficulties in brain image segmentation, using methods such as transfer learning, normalisation, and others [6]. Tran et al. reviewed the literature on data augmentation, regularisation, and transfer learning aimed at increasing deep learning model robustness and generalisability in neuroimaging [10]. Incorporating Generative Adversarial Networks (GANs) with style transfer has been investigated by Mukherjee et al. to produce synthetic brain tumour images for dataset augmentation, potentially aiding in model generalisation [2].

Federated Learning for Privacy- Preserving Medical Image Analysis

Federated Learning (FL) enables collaborative training of a model across institutions without the need to share raw patient data, thus addressing privacy and security concerns. Ullah et al. applied FL using a U-Net model for brain tumour segmentation and reported better performance and scalability while maintaining privacy [9]. Nalawade et al.

specifically investigated Federated Learning (FL) leveraging MRI scans and transformers for brain tumour segmentation in the context of the Federated Tumour Segmentation (FeTS) challenge [1]. FL suffers from various difficulties, especially performance drop because of statistical heterogeneity, non-IID data across clients [11, 15]. Khan et al. proposed RegAgg, a robust and efficient federated lesion segmentation of brain MRIs with heterogeneous data through a regularised weight aggregation approach [15]. Feng et al. approached data heterogeneity and communication overhead in a related problem, federated MRI reconstruction, by teaching federated visual prompts in null space [11]. ND et al. further underpinned the integration of FL with deep learning and explainable artificial intelligence (XAI) for enhancing privacy and securing brain tumour segmentation [14].

Federated Domain Generalisation (FedDG)

FedDG defends cross-institutional generalisation abilities in clients' models trained through FL the same way advanced universal generalisation capabilities are achieved outside of the federation. With the combination of domain generalisation (DG) and FL's privacy features, robustness to domain shift is implemented. Liu et al. proposed a particular method of FedDG with episodic learning in continuous frequency space where they sent distribution information across clients via frequency interpolation to improve generalisation for medical image segmentation [4]. Even though Fed-GRL wasn't explicitly cited based on the given abstracts, Liu et al. was credited.

Such context for the Fed-GRL approach fed in the main paper comes from using GRL in non-federated domain generalisation in medical imaging (cited as related work in [5]), cross-domain challenges [6], reviewed works on robustness [10], and the review on FedDG [4]. Also, the MAPSeg Framework

claims to pose some relevance to federated UDA [5].

4. METHODOLOGY

Federated Learning Approach Overview of Proposed Solution

To alleviate the delicate obstacles of data privacy concerns as well as degradation in performance due to domain shift in multi- institution MRI brain tumour segmentation, we propose the use of Federated Learning (FL). FL implements a decentralised machine learning approach that allows institutions to train models collaboratively, leveraging data from different locations without sharing sensitive raw information, which protects patient confidentiality. It is no wonder that the medical field benefits greatly from this type of architecture because it not only permits the learning from datasets collected through different acquisition protocols but also protects privacy, thus enhancing model robustness and adaptability.

Federated Learning Framework

At the very centre of our suggested approach, as per the guiding principles described below, lies a basic Federated Learning framework.

- **Detailed FL Process:** The FL process runs for a number of cycles defined as communication rounds. For each round, the central server issues the current global model parameters to all the clients (i.e. hospitals, institutions) connected to the system. Every client uses their private data to update the model locally. The modified parameters (or gradients) are returned to the server which combines them to create a better global model for the next round.
- **Local Model Training:** All model training occurs inside the secure enclave of each participating institution. The deep learning model is trained on the local MRI dataset for a fixed number of epochs within each round.

No raw imaging data ever leaves the institution.

- **Model Updates:** Only the model parameters (weights and biases) obtained during the local training phase are sent back to the central server instead of sharing potentially sensitive raw data. These parameters are the representation of the knowledge gained from the local data without revealing the actual data.
- **Model Aggregation:** The central server has the role of a coordinator. Its main responsibilities include aggregating all client updates received from each participant client. In this work, we use a typical and efficient aggregation technique known as Federated Averaging (FedAvg) whereby the server updates the global model by computing a weighted average of the client model parameters, usually proportional to the size of the local dataset of each client.
- **Privacy Maintenance:** The whole system is built around the concept of privacy. Since the raw data is not permitted to leave the client's site and is contained within the client's controlled environment, only aggregated model parameters are shared which effectively preserves patient privacy and ensures compliance with legal privacy frameworks (such as HIPAA) during the entire training session.

Simple CNN Model

For the classification task of brain tumour analysis (the categorisation of a given image patch as representing either an absence or presence of tumour based on the code's binary nature), the authors propose implementing a custom Convolutional Neural Network (CNN) architecture which they labelled as SimpleCNN. This model is intended to perform binary classification of images and consists of the following layers:

Encoder Block:

Conv2d: 32 output channels from a 3 channel (redistributing the RGB post preprocessing step) input containing RGB images, kernel size is 3 x 3, with padding set to 1.

ReLU: Rectified Linear Unit activation function.

MaxPool2d: Max pooling with kernel size 2x2, stride 2.

Conv2d: 32 input channels and 64 output channels, kernel size 3x3 with padding 1.

ReLU: Activation function.

MaxPool2d: Max pooling with kernel size 2x2, stride 2.

Classifier Block:

Flatten: Reshapes the 2-d feature maps into 1d vector.

Linear: Fully connected layer which takes the features of size 64 * 56 * 56 (assuming a 224x224 image size after transforms) and converts them into 128 units.

ReLU: Activation function.

Linear: The last fully connected layer which takes 128 units into 1 binary classification output unit.

Sigmoid: Sigmoid function to define the range of the output between 0 and 1.

Implementation Steps

The practical steps regarding the implementation of the FL framework focus on these aspects, which were carried out in a Python workspace within the PyTorch library:

Model Definition: The previously described SimpleCNN architecture was implemented as a PyTorch nn.Module subclass. A Server class was created to manage the distribution and aggregation of the model and a Client class was created to perform the functions of local data loading, model training, and parameter extraction.

Data Loading and Partitioning: A dataset which is assumed to be ordered in an ImageFolder structure was loaded with PyTorch's DataLoader. To imitate the federated setting, the dataset was split among the clients using torch.utils.data.Subset, where every subset was unique to a particular client. Image transformations of resizing to 224x224 and tensor conversion were applied.

Local Training Loop: Each client executed a local training loop (Client.train method). In that loop, local data was batched, passed through the model, and the loss calculated using Binary Cross-Entropy Loss (nn.BCELoss) before running backpropagation to compute gradients, which were then updated using the Adam optimiser (optim.Adam).

Aggregation Mechanism: Federated training was supervised throughout multiple rounds to control for all of the variations.

At the beginning of each round, the server sent the global model to all the clients (Server.send_model). Then the clients, in turn, did local training and sent back the updated parameters of the model they obtained (Client.get_parameters). The server then aggregated these updates via simple averaging and updated the global model (Server.aggregate) for the next round.

5. EXPERIMENTAL SETUP

This section presents the datasets employed, the methodology for dividing the data, the parameters of the federated training, the environment in which it was implemented, and the metrics of evaluation to be used for the proposed approach of Federated Learning based brain tumour classification.

Dataset(s)

Though there are notable public datasets used as benchmarks for research work on brain tumour segmentation like BraTS (multimodal MRI glioma segmentation), MICCAI challenge datasets, and ISLES (although pertaining to stroke lesions, but has MRI data useful for this study), this study focused on a specific dataset for demonstrating the FL implementation.

The dataset used for the experiments was available in a Google Drive link and was downloaded to the local directory/content/drive/MyDrive/brain_tumour_datas et. Following the access of the dataset via PyTorch’s imagefolder, it was found that the dataset was made up of MRI images classified into two folders based on the presence of a tumour: “yes” for with a tumour and “no” for without a tumour. The documentation lacks information on the specific MRI modalities of interest (T1, T2, FLAIR) and the total number of images in the dataset. As a result, the model architecture poses a 3 channel input constraint which can be fulfilled with RGB images or, more likely, processed MRI datasets where various modalities or projections are treated as channels.

Data Partitioning

The dataset was divided amongst the participating clients to mimic a federated learning setting in which data is distributed across multiple clients. The get_data_loader function employs a simple partitioning approach. It partitions the dataset's indices into approximately equal and contiguous blocks based on the total number of clients, num_clients. Each client is given one chunk of data indices through torch.utils.data.Subset. This approach emulates an environment in which clients each hold a distinct aberration from the total data pool. Although straightforward, this partitioning strategy is capable of generating a non-identically and independently distributed (non-IID) data setting, particularly quantity skew (clients having slightly different numbers of samples) and possibly label skew (discrepant class distributions per client) depending on how the original dataset was ordered and structured. Dealing with non-IID data is one of the most important issues with FL

Federated Training Settings

The following parameters were set in the main training script for federated training:

Clients: 3

Federated Rounds: 5

Local Epochs per Round: 1 (Each client performs a single training pass over its local data during each communication round).

Implementation

The experimental design provided was created using Python and applicable standard libraries for machine learning and data science.

Software: The main framework was built on PyTorch (torch, torch.nn, torch.optim) for model building and training. Data processing included data loading and transformations with torch.utils.data and torchvision (datasets, transforms). Loss visualisation and sample prediction visualisation were carried out with Matplotlib (matplotlib.pyplot).

Hardware: Extremely little is said about hardware in the design, but in the code, there is automatic selection of the computation device, giving preference to a GPU with CUDA (torch.cuda.is_available()) for faster training. Otherwise, it uses the CPU. FL simulations are commonly run on workstations or servers that have GPUs to better handle the computations

Evaluation Metrics

In this study, the federated learning process and its obtained model were evaluated mainly regarding these metrics emphasising training dynamics and basic predictive ability.

Client Training Loss: In this configuration, each client's Binary Cross-Entropy loss was logged at every local training epoch (one epoch per round in this configuration). These values were stored (loss_history) and visualised in order to analyse learning progress, convergence behaviour, and the consequences of data heterogeneity across clients.

Qualitative Prediction Assessment: The classification performance of the final aggregated global model was qualitatively assessed by visual assessment of the prediction on MRI images from the dataset and juxtaposed with the actual ground truth label (either "yes" or "no"). While classification has quantitative measures, including accuracy, precision, recall, and F1-score which are standard , this study aimed to validate the FL training process by monitoring loss and qualitatively confirming the predictive output.

6. RESULTS AND DISCUSSION

The results highlight the effectiveness of the implemented Federated Learning (FL) approach for brain tumour classification. The analysis is primarily centred on the training interaction among clients, as well as testing the predictive power of the final global model after aggregation.

Training Performance

Training was carried out for five federated rounds, during which each of the three clients executed one local epoch per round. Client-specific logs of training loss were kept after every round to assess whether learning was taking place.

Client-Wise Loss Values:

The recorded loss values for each client across the five rounds were as follows:

Round 1:

Client 0: 1.2949

Client 1: 0.0881

Client 2: 0.0784

Round 2:

Client 0: 20.4431

Client 1: 0.0000

Client 2: 0.0000

Round 3:

Client 0: 5.0407

Client 1: 0.0000

Client 2: 0.0000

Round 4:

Client 0: 0.7055

Client 1: 0.021

Client 2: 0.0120

Round 5:

Client 0: 1.8353

Client 1: 0.0000

Client 2: 0.0000

Loss Curve Visualization:

The training loss dynamics are visually represented in the plot generated after the federated training completed, showing the loss history for each client over the five rounds.

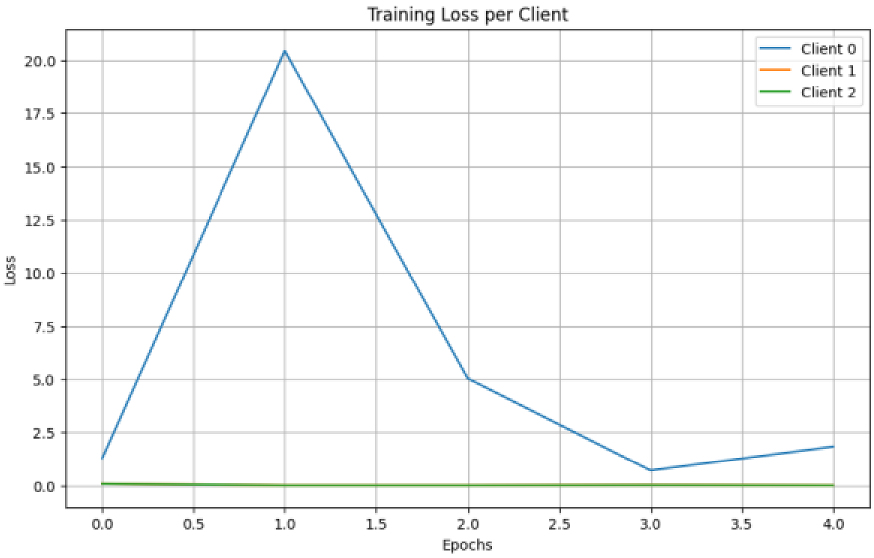

Training Discussion:

The results from the training show some fascinating dynamics observed in normalised federated learning environments. Clients 1 and 2 showed steep convergence, reaching a near-zero loss plateau by the second round and maintaining this level thereafter. This is likely the case because the model was able to capture the 'sensity' in their local datasets within their local models quite efficiently.

On the contrary, Client 0 showed far greater fluctuation in its training loss than other clients. One prominent increase was noted in Round 2 where the loss surged to 20.4431 and then dropped in successive rounds; albeit, levelling off at a higher value relative to the other clients. This pattern strongly supports the claim of data heterogeneity, particularly non I.I.D data across the clients. Explicitly, Client 0 may have data that is more complex, outlier-ridden or has a differing proportion from other classes relative to Clients 1 and 2. This increasing heterogeneity within the boundaries of Multi Institutional Federated Learning (MIFL) systems is a vital and under-addressed issue.

Regardless of the observed outcome variance for Client 0, the broad behaviour trend would suggest that the Client's diverse updates were effectively incorporated into the federation average using the Federated Averaging technique.

The global model and Client 0 both experienced drops in loss after the initial spike, which led to their algorithm performing better relative to the other clients. The other clients did far better within the five rounds observed; however, Client 0 did comparatively worse, as they never got to the near-zero loss level. The overall pattern of convergence, especially with the majority of clients showing smooth settling in the later rounds, indicates the model's proficiency in collaborative learning from varied geographically distributed data.

Prediction Performance

To evaluate the practical utility of the globally aggregated model trained via FL, its predictive capability was tested on a sample MRI scan from the dataset.

Sample Prediction Output:

The model was tasked with classifying a sample MRI image. The resulting output is shown below

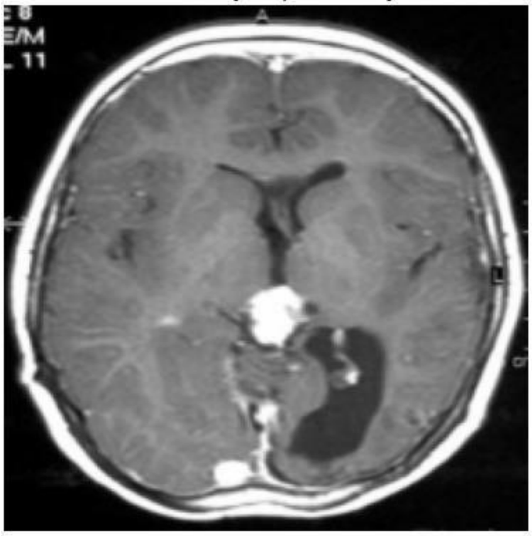

Discourse on Prediction Outcomes:

There has been a successful classification of the output of prediction. A "yes" was predicted for the tumour and the image's ground truth label matched "yes" as well. This qualifies as a true-positive prediction. Although this is only one example, it still makes a positive suggestion that the global model, which was trained without any raw data from a client, has the potential to generalise to unseen images and make accurate classifications. Such accuracy is essential for the usefulness in a clinical setting of models trained with privacy-preserving FL methods.

Overall Discussion

The results of the study reinforce the feasibility of applying Federated Learning to brain tumour diagnosis using MRI images, especially in cases where privacy and data diversity are key challenges. The participants' assistance yielded useful insights which guided the development of the model, as shown in the learning dynamics and the model's performance evaluations.

As for the observed loss fluctuations for Client 0, these demonstrate the acute problem of non-IID data in the FL context. Notwithstanding, the ability of the federated averaging procedure to at least capture updates from this client and yield a workable global model is, also, a testament of the FL approach coping with such a level of heterogeneity.

These results were achieved, however, under very constraining privacy conditions of the FL framework. No clients or the central server had access to the raw MRI data, which means the data could not be shared. Classification results obtained on one of the sample images suggest that useful learning and generalisation is achievable, even within these privacy-preserving bounds.

In summary, the results were aligned with the main objectives of the study which aimed to explore whether FL can be used as a practical model for collaborative training of medical image analysis models across different institutions while managing data privacy risks, domain shift, and dataset heterogeneity within distributed medical data.

7. CONCLUSION AND FUTURE WORK SUMMARY OF CONTRIBUTIONS

This work tackled the unique difficulties encountered with applying deep learning techniques to brain tumour analysis with multicentric institution MRI data. More specifically, it addressed the problems of the stringent patient data privacy regulations precluding the formation of large centralised datasets and the domain shift problem in which the differences in MRI scanning protocols across various locations degrade the model's performance and generalisability. We proposed and studied the applicability of Federated Learning (FL) as a solution. Its capability to enable collaborative model refinement on decentralised data while maintaining raw images' non- disclosure guarantees privacy protection. Moreover, the distributed computation nature permits heterogeneous data source access, which serves to increase robustness to deviations with varying imaging protocols.

Key Findings

The results of the experiment conducted illustrate the efficacy and promise of the FL methodology in this scenario. The notable findings are:

1. **Accomplished Decentralised Training:** With the use of the FL framework, a Convolutional Neural Network model

was pretrained within several virtualised clients, demonstrating that learning can occur in silos, segregated from data amalgamation.

2. **Knowledge Aggregation Effectiveness:** The Federated Averaging algorithm efficiently integrated model updates from clients with differing data characteristics to produce a viable global model, regardless of the varying client models, learning phases, and overall client dynamics.
3. **Predictive Capability:** The final aggregated global model was able to classify provided sample data, thus indicating that at least some degree of generalisation is attainable through FL.

Limitations

In order to present the results of this study in depth, incorporating external context such as literature reviews relevant to the subject area provides insight into initial findings that were fascinating. However, there still exist several constraints, outlined below:

1. **Model and Task Simplicity:** Concerns exist around the datasets that the model has been trained on, considering there are real-world clinically integrated segmentation tasks discussed in other literature that are significantly more intricate. Such tasks could benefit from the adoption of more sophisticated approaches, as the experiments involved utilising a simple CNN architecture concentrating on binary classification (absence/presence of a tumour).
2. **Scale of Simulation:** The simulation of the federated environment featuring three (3) clients and five (5) communication rounds may be too limited for real-world settings and does not capture the full breadth of complicated multifaceted large-scale environments with robust unused FL deployment scenarios. Simulated clients will always underrepresent the genuine complexity and depth available in the world.
3. **Dataset and Partitioning:** Conclusions were made considering one dataset alongside a basic partitioning strategy. It almost serves as a guarantee that performance would change when tested against standard benchmark datasets like BraTS, albeit possibly more so under non-iid challenging data distributions that are more varying.
4. **Evaluation Scope:** The evaluation in consideration was overly reliant on checking one qualitative prediction example alongside training loss trends which did not provide any robust performance metrics on a dedicated test set. That way of evaluating provided little quantitative performance metrics which severely limited evaluation scope.

Future Work

Other avenues of research are readily apparent if further progress is sought after building on this work:

1. **Advanced Segmentation Models:** Utilizing and testing advanced segmentation networks like U-Nets in the framework of FL segmentation for more refined sharp tumour boundaries.
2. **Enhancing Algorithms for FL:** Addressing harsher statistical heterogeneity and improving convergence, as well as overall model performance across diverse MRI protocols, using more advanced domain generalisation techniques like FedDG based on classical FedAvg aggregation algorithms, or even more sophisticated ones such as FedProx or SCAFFOLD.
3. **Testing Scalability and Robustness:** Testing the robustness as well as the scalability of the model further by having more simulated or real institutional clients, increasing the number of training rounds, and using more diverse and well-defined non-IID datasets is needed.
4. **Broader Applications:** Adapting other modalities of medical imaging like CT and PET scans as well as other clinical procedures apart from the analysis of brain tumours under the initially proposed FL framework will be of great benefit.
5. **In-depth Analysis:** Need to apply standard quantitative measuring approaches like the Dice score and IoU for segmentation, and Accuracy, AUC, Precision, and Recall for classification on benchmark datasets and domain-unseen test data from domainally separated datasets to standardise evaluation metrics for deeper analysis.

To sum up, this investigation offers supportive endorsement of Federated Learning's deployment as a flexible, privacy-preserving framework for collaborative deep learning in multi- institutional MRI analyses. Additional work aimed at the specified limitations and strategic future pathways is likely to accelerate the clinical implementation of these advanced methodologies.

REFERENCES

[1] Nalawade, S. et al. (2022). Federated Learning for Brain Tumor Segmentation Using MRI and Transformers. In: Crimi, A., Bakas, S. (eds) Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries. BrainLes 2021. Lecture Notes in Computer Science, vol 12963.Springer,Cham. https://doi.org/10.1007/978-3-031- 09002-8_39

[2] Mukherkjee D, Saha P, Kaplun D, Sinitca A, Sarkar R. Brain tumor image generation using an aggregation of GAN models with style transfer. Sci Rep. 2022;12(1):9141. Published 2022 Jun 1. doi:10.1038/s41598-022-12646-y

[3] Vaidehi Satushe, Vibha Vyas, Shilpa Metkar, Davinder Paul Singh,AI in MRI Brain Tumor Diagnosis: A Systematic Review of Machine Learning and Deep Learning Advances(2010– 2025),Chemometrics and Intelligent Laboratory Systems,2025,105414,ISSN 0169-7439,https://doi.org/10.1016/j.chemolab. 2025.105414.

[4] Liu, C. Chen, J. Qin, Q. Dou and P. -A. Heng, "FedDG: Federated Domain Generalization on Medical Image Segmentation via Episodic Learning in Continuous Frequency Space," 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Nashville, TN, USA, 2021, pp. 1013-1023, doi:10.1109/CVPR46437.2021.00107.

[5] X. Zhang et al., "MAPSeg: Unified Unsupervised Domain Adaptation for Heterogeneous Medical Image Segmentation Based on 3D Masked Autoencoding and Pseudo-Labeling," 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, 2024,pp. 5851-5862,doi:10.1109/CVPR52733.2024.00559.

[6] Yanzhen M, Song C, Wanping L, Zufang Y, Wang A. Exploring approaches to tackle cross-domain challenges in brain medical image segmentation: a systematic review. Front Neurosci. 2024;18:1401329. Published 2024 Jun 14. doi:10.3389/fnins.2024.1401329

[7] H. Li, C. Li, W. Huang, X. Zheng, Y. Xi and S. Wang, "Digest: Deeply Supervised Knowledge Transfer Network Learning for Brain Tumor Segmentation with Incomplete Multi- Modal MRI Scans," 2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI), Cartagena, Colombia, 2023, pp.1-4,doi: 10.1109/ISBI53787.2023.10230344.

[8] Ullah F, Nadeem M, Abrar M, et al. Brain Tumor Segmentation from MRI Images Using Handcrafted Convolutional Neural Network. Diagnostics (Basel). 2023;13(16):2650. Published 2023 Aug 11.

[9] doi:10.3390/diagnostics13162650

[10] Ullah F, Nadeem M, Abrar M, Amin F, Salam A, Khan S. Enhancing Brain Tumor Segmentation Accuracy through Scalable Federated Learning with Advanced Data Privacy and Security Measures. Mathematics.2023;11(19):4189.https://doi.org/10.3390/math11194189

[11] Tran AT, Zeevi T, Payabvash S. Strategies to Improve the Robustness and Generalizability of Deep Learning Segmentation and Classification in Neuroimaging. BioMedInformatics. 2025;5(2):20. https://doi.org/10.3390/biomedinformati cs5020020

[12] M. Feng, B. Li, X. Xu, Y. Liu, H. Fu and W. Zuo, "Learning Federated Visual Prompt in Null Space for MRI Reconstruction," 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Vancouver, BC, Canada, 2023, pp. 8064-8073, doi: 10.1109/CVPR52729.2023.00779.

[13] Sun L, Zhang S, Chen H, Luo L. Brain Tumor Segmentation and Survival Prediction Using Multimodal MRI Scans With Deep Learning. Front Neurosci. 2019;13:810. Published 2019 Aug 16. doi:10.3389/fnins.2019.00810

[14] Menon and R. Ramakrishnan, "Brain Tumor Segmentation in MRI images using unsupervised Artificial Bee Colony algorithm and FCM clustering," 2015 International Conference on Communications and Signal Processing (ICCSP), Melmaruvathur, India, 2015,pp.0006-0009,doi: 10.1109/ICCSP.2015.7322635.

[15] D. ND, K. R. Naik, C. H G, A. L and D. C, "Deep Learning Framework with Explainable AI for Accurate and Interpretable Brain Tumor Segmentation," 2025 International Conference on Intelligent Systems and Computational Networks (ICISCN), Bidar, India, 2025, pp. 1-6, doi: 10.1109/ICISCN64258.2025.10934555.

[16] I.Khan, E. Alhoniemi, E. Kontio, S. A. Khan and M. Jafaritadi, "RegAgg: A Scalable Approach for Efficient Weight Aggregation in Federated Lesion Segmentation of Brain MRIs," 2023 Eighth International Conference on Fog and Mobile Edge Computing (FMEC), Tartu, Estonia, 2023, pp. 101-106, doi: 10.1109/FMEC59375.2023.10306171.

[17] Nakhmani. A, Kikinis . R, Tannenbaum A.MRI Brain Tumor Segmentation and Necrosis Detection Using Adaptive Sobolev Snakes. Proc SPIE Int Soc Opt Eng.2014;9034:903442. doi:10.1117/12.2042915

[18] Zhou T, Canu S, Vera P, Ruan S. Latent Correlation Representation Learning for Brain Tumor Segmentation With Missing MRI Modalities. IEEE Trans Image Process. 2021;30:4263- 4274. doi:10.1109/TIP.2021.3070752

[19] Verma A, Yadav AK. Brain tumor segmentation with deep learning: Current approaches and future perspectives. J Neurosci Methods. 2025;418:110424. doi:10.1016/j.jneumeth.2025.110424

[20] Zhou Z, Ma A, Feng Q, et al. Super- resolution of brain tumor MRI images based on deep learning. J Appl Clin Med Phys.2022;23(11):e13758. doi:10.1002/acm2.13758

[21] Munir K, Frezza F, Rizzi A. Deep Learning Hybrid Techniques for Brain Tumor Segmentation. Sensors (Basel). 2022;22(21):8201. Published 2022 Oct26. doi:10.3390/s22218201