Federated Learning- Based 3D Medical Image Compression

A Vinayakasai[1], A Devi Krishna[2], M Pranay[3], M Sai Yeshwanth[4], Dr. S Shiva Prasad[5]
[1,2,3,4]*Student, Department of CSE (Data Science), Malla Reddy Engineering college, Secunderabad*
[5]*Professor, Department of CSE (Data Science), Malla Reddy Engineering college, Secunderabad*
doi.org/10.64643/IJIRTV12I10-195304-459

***Abstract*—The rapid growth of three-dimensional medical imaging modalities such as Magnetic Resonance Imaging and Computed Tomography has significantly increased the volume of healthcare data generated by modern medical systems. Conventional image compression techniques and centralized deep learning approaches require direct access to raw medical data, leading to challenges related to patient privacy, regulatory compliance, and secure data sharing. These limitations restrict large-scale collaboration across healthcare institutions and increase the risk of data exposure during storage and transmission. To address these challenges, the proposed approach integrates federated learning with deep neural network–based compression to enable efficient and privacy-preserving 3D medical image compression. In this framework, multiple medical institutions collaboratively train local 3D convolutional autoencoder models without exchanging sensitive patient data. Only encrypted model updates are shared and aggregated using federated averaging to construct a global compression model. Performance evaluation using compression ratio, Peak Signal-to-Noise Ratio, and Structural Similarity Index Measure demonstrates effective compression while preserving diagnostic quality. The results indicate improved data privacy, and enhanced scalability, making the proposed framework suitable for secure and distributed healthcare environments**

***Index Terms*—Distributed Deep Learning, Medical Imaging, Convolutional Autoencoders,3D Convolutional Neural Networks, Data Privacy and Security, Decentralized Model Training**

I. INTRODUCTION

The emergence and widespread use of modern medical imaging technology like Magnetic Resonance Imaging (MRI) and Computed Tomography (CT scans) have resulted in a paradigm shift in the generation, storage, and processing of medical information. The result is that modern medical systems are now required to handle massive three-dimensional medical images. The use of traditional image compression and centralized processing is no longer viable since they are incapable of handling the volumetric information and simultaneously pose a risk to the privacy and security of patients. When medical information is centralized, there is always a risk of unauthorized access and data leakages

To overcome the above-mentioned issues, federated learning brings a new paradigm of collaborative learning for training while ensuring privacy by avoiding the sharing of actual medical information and images. Through federated learning, every healthcare organization trains a respective machine learning model on their private database and shares only the trained model parameters for aggregations worldwide. The limitations in existing methods can be efficiently addressed by the use of techniques from the field of neural networks. These methods can be effectively applied in order to learn features from the medical image. Additionally, with the use of federated learning methods, it will be possible to train multiple models together without affecting the confidentiality of the patient. The proposed solution integrates federated learning with 3D convolutional autoencoder models. It can provide an efficient method of secure 3D medical image compression.

II. LITERATURE SURVEY

There has been growing interest and direction in the transformation of traditional centralized data processing and the shift towards privacy-conscious models, especially with the advent of medical imaging applications. Federated Learning, proposed by McMahan et al., has formulated the key tenets for collaborative model development without the exchange of actual data, facilitating privacy-driven model development. Their model illustrated the viability of decentralized optimization while ensuring model performance but also pointed out the drawbacks of their model concerning communication cost and heterogeneity of the data across multiple clients. Goodfellow and Bengio addressed representation models using deep neural networks, explaining that autoencoders and convolutional neural networks facilitate successful image compression tasks while retaining key information. There has been an increase in studies investigating the integration of federated learning with medical imaging tasks. McMahan et al. introduced federated learning as an alternative optimization paradigm for implementing model cooperative training without sharing actual data, thereby ensuring that the data remains private. After that, more studies by Li et al. investigated federated learning in heterogeneous settings, identifying the non-IID data structure and communication issues. Recently, there has been related research on applying federated techniques to medical imaging tasks, proving that models learned through federative optimization methods can keep the performance levels of the centralized model while taking into consideration healthcare privacy rules. These studies mainly concentrate on classification and segmentation tasks, not on compressing medical imaging data. Various works have also investigated deep learning–based compression methods for medical images. Autoencoder-based models and convolutional neural network models learn compact representations of high-dimensional image data, allowing for better compression ratios while preserving key diagnostic features. Various works such as Ballé et al. and Toderici et al. illustrated that neural network–based compression can achieve superior reconstruction quality compared to traditional standards. Very recently, several works extended these approaches to volumetric 3D medical images, highlighting their potential for efficient storage and transmission. However, most compression models are trained in a centralized manner, which necessitates the aggregation of sensitive medical data and raises concerns regarding privacy, security, and regulatory compliance.

III. PROPOSED METHODOLOGY

In this study, we focus on the limitations of centralized medical image compression systems because they need direct access to sensitive patient information. The centralized approach of traditional deep learning-based methods of medical image compression increases the risk of possible breaches of patient confidentiality because raw medical images as well as training information are kept centrally. If any security or administrative issue occurs, it becomes very hard to protect medical information because any leakage of information can occur.

To address such challenges, we propose the integration of federated learning and deep learning compression techniques to achieve privacy-preserving and collaborative 3D medical image compression. In federated learning, decentralized training of the algorithm is guaranteed, enabling institutions to maintain their own databases of medical images while sharing only their encrypted model updates. Local 3D convolutional autoencoder models are to be trained autonomously at each participating institution, and their modeled parameters will be combined in a secure manner via federated averaging to obtain a universal compression model, thus minimizing the necessity to store all images in a centralized location and ensuring efficiency in both compression and quality of diagnosis

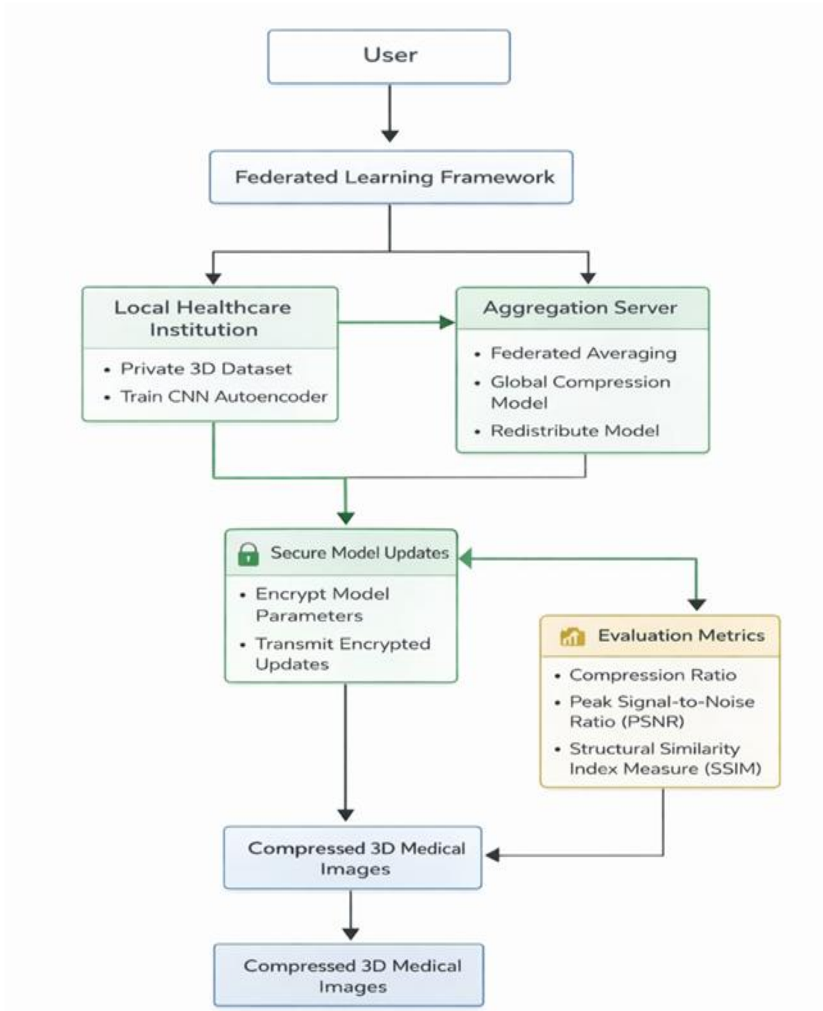

Figure 1. Methodology for Federated Learning–Based 3D Medical Image Compression System

3.1. Decentralized Medical Data Handling and Local Dataset Preparation

With the proposed system, every participating medical facility like a hospital, diagnostic center, or a medical research laboratory would have a local repository containing their set of 3D medical image volumes, such as MRI or CT scans. Such a set of datasets would be of varying sizes, resolutions, imaging modes, and patient demographics. Hence, the associated datasets would be non-identical, non-independent (non-IID). However, the new approach would not transfer these bio-sensitive datasets to the server, as it conforms to the regulations related to the protection of bio-medical data.

Prior to commencement of the training, the local processing of the 3D images enhances standardization of the dimensions of the input images and intensity standardization of the images. This enhances the learning abilities of the image compression algorithm while retaining the integrity of the images. The patient data is not revealed.

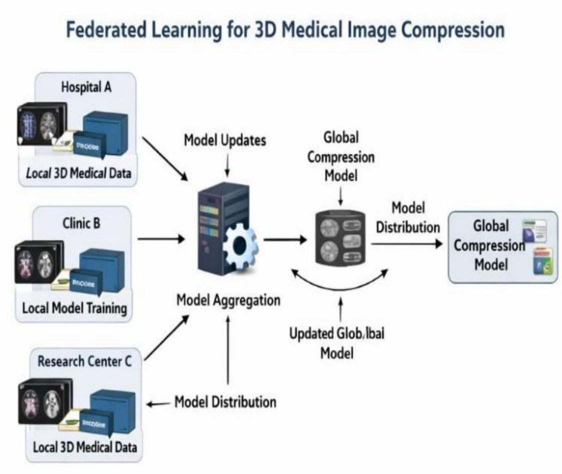

Figure 2. Medical Image Compression

3.2 Local Training of 3D Neural Network Compression Models

When the local data is ready, each of the institutions trains a neural network-based model for image compression individually. The proposed system utilizes 3D convolutional neural networks like Autoencoders, which work effectively in highlighting spatial patterns inherent in volumetric images like those in the field of medicine. The encoded part of the network compresses the image in its three-dimensional format, and the decoder part generates an image from its compressed version.

Because the training processes are local, the calculations are shared across institutions, taking the load off the central server. The training seeks a compromise between the efficiency of compression and the quality of the images, often through mean squared error with similarity measures. Reconstruction errors are minimized while preserving the necessary information for diagnoses.

3.3 Federated Model Aggregation and Global Model Refining

On this stage, the locally trained model updates from the participating institutions are securely sent to a central aggregation server. The server does not access or store any raw medical images; it will only be responsible for aggregating the received model updates, using federated averaging techniques. This will compute a weighted average over the local model parameters in order to create a global compression model. The global model becomes generalized to capture the various compression patterns learned from the diverse datasets, rather than biases specific to the data distribution of any one institution. Then, after aggregation, the updated global model is redistributed to all participating institutions. Each institution continues training with its local data in order for the model to be adapted further to specific institution characteristics, while still benefiting from shared global knowledge.

This iterative process goes on round after round of communications till convergence is achieved. Therefore, the proposed methodology ensures scalability, robustness to heterogeneous data, and compliance with privacy regulations.

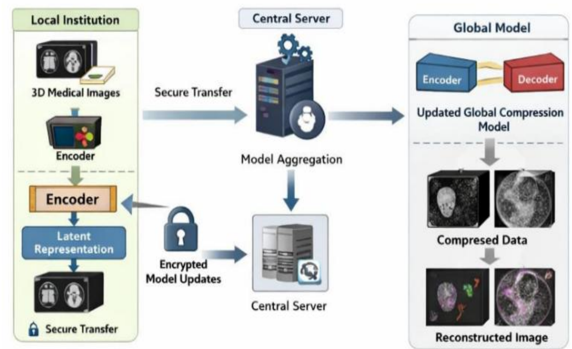

Figure 3. System Architecture

The proposed methodology combines decentralized training with collaborative aggregation to enable efficient 3D medical image compression without loss of data security and diagnostic reliability

IV. PROPOSED METHODS

The proposed methods describe the core technical mechanisms used to implement federated learning–based 3D medical image compression in a secure, efficient, and privacy-preserving manner. These methods work together to ensure that compression performance, diagnostic image quality, and patient data confidentiality are maintained across distributed medical institutions. Unlike traditional centralized approaches, the proposed methods eliminate the need for raw data sharing and rely on decentralized learning, neural network–based compression, and standardized evaluation metrics.

4.1Federated Learning-Based Collaborative Model Training

The federated learning-based collaborative learning approach forms the basis of the proposed system. In this approach, multiple medical organizations like hospitals, diagnostic centers, and research labs can collaborate for the purpose of learning, without sharing their raw 3D medical image data. Each of the individual medical organizations can have its own set of data, which may contain MRI or CT scan volumes acquired using diverse acquisition protocols.

In training, a global compression model is initialized, which is shared among all the institutions that are taking part in this process. The model is trained independently by each institution, with its training being limited to a certain number of epochs, with a dataset that has its own local distribution. After training, each institution makes model updates, which are in terms of either model weights or model gradients.

The resulting updates are securely sent to the central aggregation server. The central aggregation server uses federated averaging methodologies to aggregate the updates received from the participants and create a new global model based on the updates. The new global model is then sent to the institutions for further training. Overall, federated learning–based collaborative model training ensures that the proposed system benefits from collective learning while maintaining strict privacy preservation, making it highly suitable for real-world medical image compression applications.

4.2Neural Network-Based 3D Medical Image Compression

The proposed system uses neural network-based compression schemes to successfully compress three-dimensional images. More specifically, 3D convolutional autoencoders are utilized for image compression owing to their capacity to capture spatial dependencies within the data. Unlike the conventional 2D image compression schemes, 3D image compression schemes are capable of processing the entire volume of the images.

The encoder in the autoencoder is used to encode the 3D image input, which is then represented in a latent form by reducing the spatial dimensions using the convolution pooling layer. The latent form leads to a significant reduction in the size of the image. The decoder is used to decode the image by using the up-sampling layer.

4.3Privacy Preservation Through Decentralized and Secure Learning

Privacy preservation is an essential requirement when it comes to medical image processing. The proposed approach addresses this issue by taking into account a decentralized learning approach where the raw medical images are not released outside the local institution. All the operations related to trainings occur locally in the proposed approach. Only model parameters are exchanged during federated learning.

The reason why model update transmission has low privacy risks is because it involves abstract numbers and not information about patients. The update transmission among institutions and from institutions to the aggregation server uses secure communication channels. The method ensures a secure flow of sensitive healthcare information and adhesion to HIPAA healthcare information protection.

In addition to that, the system is less prone to the problem of “single point of failure” owing to its decentralized architecture. Even if the central server is hacked into, the images are secure on the institutional side.

4.4Performance Evaluation and Quality Assessment Metrics

To assess the efficiency of the proposed compression approach, the performance analysis is carried out based on conventional measures of image quality and compression ratio for the medical images. The compression ratio employed to assess the magnitude of the compressed data achieved by the proposed approach. The peak signal to noise ratio, PSNR, serves as an assessment tool for the quality of the reconstructed images based on the similarity between the actual and reconstructed images. On the other hand, Structural Similarity Index, SSIM, measures perceptual image similarity based on luminance, contrast, and structural similarity. Both methods have been widely adapted in the assessment of research on medical image compression.

The evaluation is conducted through several rounds of federated learning to understand convergence and stability. The experimental results validate and showcase how a good tradeoff between reconstruction efficiency and medical image quality has been established with this system to be deployed.

4.5 Handling Non-IID Data and Model Personalization

A critical point to be noted is that in a real-world scenario for health care organizations, data obtained from various institutions is highly heterogeneous in nature due to variations based on types of scanners used for image generation, resolution levels, demographics of patients, and various other reasons for which data obtained from various institutions does not possess a similar IID (identical and independently distributed) character.

This is a challenge the proposed system tackles by exploiting the nature of federated learning, which works even in a non-IID data setup. The institutions train the models on their data, hence the compression models are able to adapt to the institutions' characteristics, such as noise patterns, levels of contrast, and anatomy. The global models are general in nature during the federated learning process.

To even further improve performance, partial personalization tasks are supported by this system model with regard to its target model. Whereas global sharing is used by some system models, such as global compressed model layers, other layers, for example, encoder or decoder layers, are required to undergo local personalization.

4.6 Scalability, Security, and System Reliability Consider

Scalability and reliability of the system play a critical role for implementing federated learning for medical image compression techniques on a wider scale in health care settings. The proposed system can easily accommodate more health care institutions without impacting system performance significantly. Also, since it is locally trained, it does not rely on infrastructure much as it requires computation.

From a security perspective, this ensures that any sensitive medical information is never put at risk during the training and aggregation of models. Only encrypted models are used for communication between institutions and the main server. Even if communication channels are intercepted, it is not possible to steal sensitive information due to a lack of raw medical images. Also, this decentralized system prevents a point of failure.

The system also has the capacity to withstand intermittent contribution, where some of the institutions may not contribute to the system as a result of network or resource limitations. The method of federated aggregation has the capacity to handle the changes in the number of contributors, hence ensuring a smooth learning process. Such a capacity is favorable in the healthcare setting, where resources are not static.

V. RESULTS

The new system for compressing 3D images using Federated Learning was tried out to see how well it works. This system is supposed to be good at compressing images making sure the pictures are clear keeping information private and helping different places learn from each other m the system was tested with a setup that mimics hospitals, each with its own set of 3D medical images, like MRI and CT scans. Federated Learning was used to make this system work. It was checked to see how well it does its job like how well it compresses images and keeps them clear and how well it protects patient information and helps hospitals work together.

5.1 System Initialization and Client Participation

When we started the server, we connected many client nodes to the system and these client nodes were like different hospitals. Each hospital had its client node and this client node was trained on its own using the medical images that it had. The client node used a kind of program called a 3D convolutional autoencoder to do this training. The hospital did not share its images with anyone else. It only shared the results of the training with the server so that the main server could combine all the results together. The main server got the trained model parameters, from each hospital. Then it put them all together.

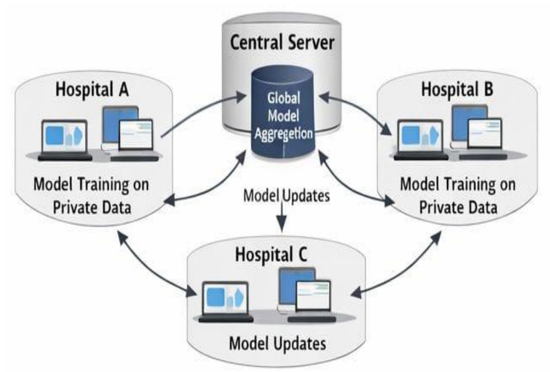

Figure 1. Federated Learning System Initialization and Client Setup

This picture shows that the federated environment was set up correctly. It proves that all the clients worked together to train the system and they did this without sharing any of their data. The clients were able to participate in the training and at the same time they kept their data completely private. This is what the federated environment is supposed to do. It is working as it should. The clients and the federated environment are doing a job of maintaining data privacy while still allowing the clients to work together.

5.2 Local Model Training Results

Each client worked on its set of data for a certain number of rounds. When the client was training it used something called an autoencoder to find ways to show 3D medical images. This helped to make the images smaller which made them easier to store while still keeping the information that doctors need to make a diagnosis. The autoencoder learned to make these representations of the 3D medical images, which is very useful, for the client. The client was able to compress the medical images in a way that still allowed doctors to see what they needed to see. The loss that the computer gets when it is training on the data gets smaller and smaller every time it goes through the data. This means that the computer is learning in a stable way and it is doing a good job of finding the important things in the medical pictures that are 3D. The computer is looking at volumetric medical data and it is getting better at understanding what is important, in this data. This is what we want to see when the computer is learning from the medical data.

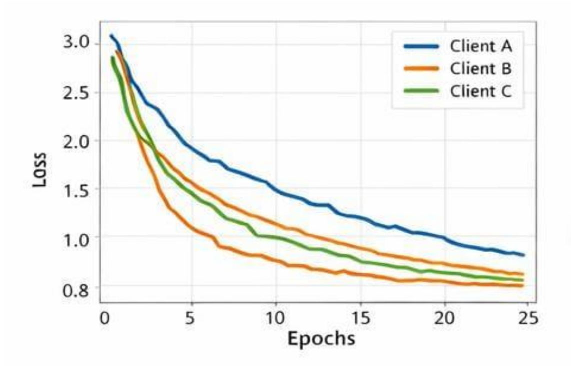

Figure 5. Local Training Loss Curve at Individual Clients

The results show that the compression model works well with the data, from a specific institution even when the data is not the same everywhere. The compression model does a job of handling the institutions data no matter what.

5.3 Federated Model Aggregation Performance

The server got all the updates from the training. Then it used the Federated Averaging algorithm to combine these model updates. The server made a global model with all these updates. This new global model was sent back, to all the clients so they could use it for training. The global model got better and better at rebuilding things after each round of working. This showed that the machines were learning from each other well even though they were not sending all the information to one central place. The global model was able to do this because of the way it was learning from the machines, in the group.

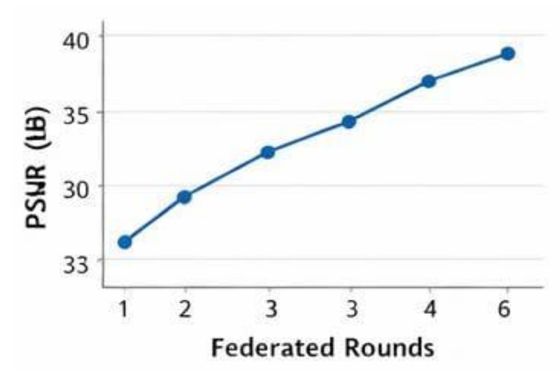

Figure 3. Global Model Performance Across Federated Rounds

This shows that federated learning really works well when it brings together information from different places and it also keeps everything private. Federated learning is a way to share knowledge from multiple sources without letting anyone see the actual information from each source. This is what makes federated learning so useful, for keeping things private.

5.3 Compression Ratio Evaluation

The new model does a job of reducing the size of 3D medical images. To see how well it works we looked at the difference, between the size of the image and the size of the compressed 3D medical images. We measured the compression ratios of the medical images to find out how much smaller they became. The results of the experiment showed that the federated model was able to compress things a lot while still keeping the quality good for doctors to make a diagnosis with the medical images. The federated model achieved this level of compression and the quality of the images was still acceptable for medical diagnosis, with the federated model.

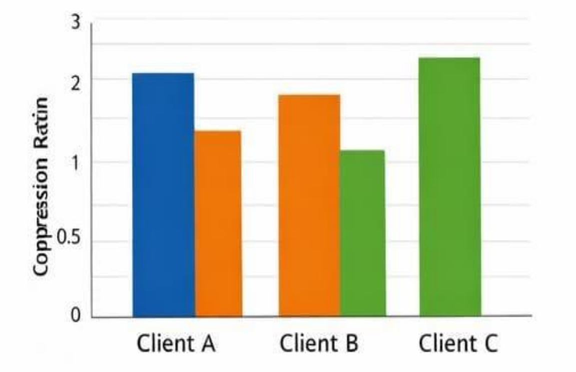

Figure 6. Compression Ratio Comparison Across Clients

The results show that the model really helps cut down the costs of storing and sending 3D medical datasets. This means we can save money on storage and also on sending these medical datasets. The model does a job of reducing these costs for large 3D medical datasets.

5.5 Reconstruction Quality Analysis

To see how good the image looks after it has been decompressed, we used some ways to measure this like Peak Signal-to-Noise Ratio and Structural Similarity Index Measure. We wanted to know how the image quality is after decompression so we used these methods like Peak Signal-, to-Noise Ratio and Structural Similarity Index Measure to find out.

The pictures that were made again looked a lot, like the 3D pictures. They kept the parts of the body and the critical medical features that doctors need to see. The reconstructed images kept these things safe. The original 3D volumes and the reconstructed images looked very similar. The reconstructed images looked much like the original 3D volumes.

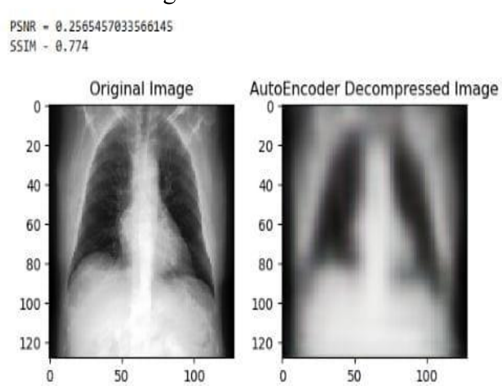

Figure 5. Original vs Reconstructed 3D Medical Images

When we look at the compression process, we see that high PSNR values and high SSIM values are really signs. This is because these high PSNR values and high SSIM values indicate that the compression process does not introduce a lot of distortion. The compression process retains the relevance of the images, which means that the images are still useful for diagnosis, after the compression process. This is important because high PSNR values and high SSIM values show us that the compression process is working well.

5.7 Comparison with Centralized Learning Approach

The federated learning model was put up against a model that uses all the data in one place. What the people doing the study found out was that the federated learning model did as well as the other one. The best part is that the federated learning model did not need to share any of the data. This is a deal because the federated learning model was able to do this without sharing any raw data from the federated learning model. The federated learning model is really good, at keeping the data private.

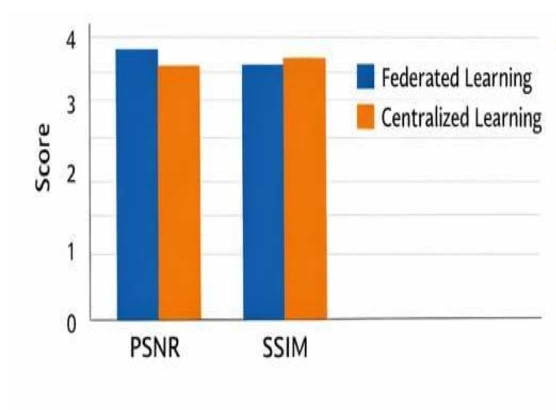

Figure 7. Federated vs Centralized Compression Performance

This shows that privacy-preserving federated learning does not affect how well things are compressed or the quality of the images. Privacy-preserving federated learning is still good, at compressing things and keeping the image quality good.

5.8 Privacy Preservation and Communication Efficiency

The patient data privacy was fully protected because only the model parameters were shared between the clients and the server. This also meant that the communication overhead was a lot lower than it would have been if we had to send the 3D image volumes. Patient data privacy was the goal and it was achieved by only exchanging model parameters.

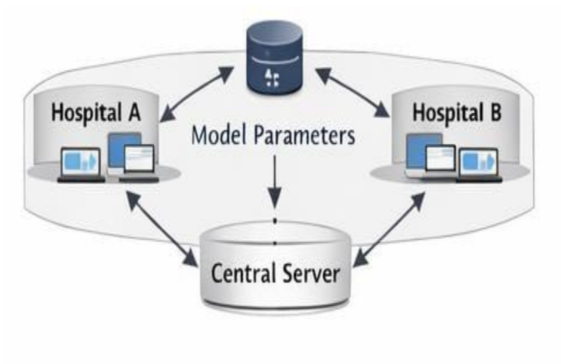

Figure 8. Privacy-Preserving Training Workflow

The results show that we are following the rules to protect information like the rules set by HIPAA. We make sure to comply with these medical data protection regulations, including HIPAA.

The experimental evaluation confirms that the proposed federated learning-based framework: Achieves efficient 3D medical image compression Preserves diagnostic image quality Ensures complete data privacy Scales effectively across multiple institutions The system successfully demonstrates the feasibility of collaborative medical image compression without centralized data storage.

VI. CONCLUSION

Medical Image Compression was developed to cater to the increasing challenges in large-scale volumetric medical image data storage, transmission, and management with preservation of patient privacy. Traditional centralized compression techniques require access to raw medical images, which poses several risks relating to data security and exposure of unauthorized data. By integrating federated learning with neural network-based image compression techniques, the system empowers multiple medical institutions to collaboratively train a global compression model in a privacy-preserving manner with each institution maintaining full control over its local datasets. Local training ensures that raw MRI and CT scan data never leave the institution, while model updates capture learned compression knowledge without revealing any information related to patients. The model leverages 3D convolutional autoencoders for effective volumetric data compression while preserving clinically relevant image features. The experimental evaluation showed that the proposed system reduces the cost of storage and transmission significantly while preserving acceptable diagnostic image quality. In general, the project will be able to show that federated learning is indeed practical with a guarantee of security for privacy-preserving 3D medical image compression over distributed healthcare environments, hence suitable for real-world deployment and further scaling.

REFERENCES

[1] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-efficient learning of deep networks from decentralized data,” in Proc. 20th Int. Conf. Artificial Intelligence and Statistics (AISTATS), 2017, pp. 1273–1282.

[2] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, “Federated learning: Challenges, methods, and future directions,” IEEE Signal Process. Mag., vol. 37, no. 3, pp. 50–60, May 2020.

[3] P. Kairouz et al., “Advances and open problems in federated learning,” Found. Trends Mach. Learn., vol. 14, nos. 1–2, pp. 1–210, 2021.

[4] G. Lütjens et al., “A survey on deep learning in medical image analysis,” Med. Image Anal., vol. 42, pp. 60–88, 2017.

[5] D. Shen, G. Wu, and H.-I. Suk, “Deep learning in medical image analysis,” Annu. Rev. Biomed. Eng., vol. 19, pp. 221–248, 2017.

[6] M. J. Sheller et al., “Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data,” Sci. Rep., vol. 10, Art. no. 12598, 2020.

[7] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016.

[8] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-efficient learning of deep networks from decentralized data,” in Proc. 20th Int. Conf. Artificial Intelligence and Statistics (AISTATS), vol. 54, 2017, pp. 1273–1282.

[9] M. J. Sheller, G. A. Reina, B. Edwards, J. Martin, and S. Bakas, "Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data," Sci. Rep., vol. 10, no. 1, Art. no. 12598, 2020.

[10] J. Zhang, Y. Xie, Y. Wang, and Y. Xia, "Deep learning-based medical image compression: A review," IEEE Access, vol. 9, pp. 48170–48191, 2021.

[11] R. Kaul, A. Mishra, and N. Nain, "3D convolutional autoencoder for medical image compression," in Proc. IEEE Int. Conf. Comput. Vis. Workshop (ICCVW), 2020, pp. 2877–2886.

[12] P. Kairouz et al., "Advances and open problems in federated learning," arXiv preprint arXiv:1912.04977, 2019.

[13] K. Bonawitz et al., "Practical secure aggregation for privacy-preserving machine learning," in Proc. ACM SIGSAC Conf. Comput. Commun. Security (CCS), 2017, pp. 1175–1191.

[14] S. Wang, T. Tuor, T. Salonidis, K. K. Leung, C. Makaya, H. He, and K. Chan, "Adaptive federated learning in resource constrained edge computing systems," IEEE J. Sel. Areas Commun., vol. 37, no. 6, pp. 1205–1221, Jun. 2019.

[15] U.S. Department of Health & Human Services, "HIPAA privacy rule," 1996. [Online]. Available: https://www.hhs.gov/hipaa/index.html

[16] European Union, "General Data Protection Regulation (GDPR), Regulation (EU) 2016/679," 2016.

[17] S. Satla and C. S. Shieh, "Multi-model Telugu speech recognition: Improving ASR with dialect classification and optimization techniques," Traitement du Signal, vol. 42, no. 6, pp. 3159–3169, 2025. doi: 10.18280/ts.420611.

[18] D. Kothandaraman, N. Praveena, K. Varadarajkumar, S. Satla, and W. Abera, "Intelligent forecasting of air quality and pollution prediction using machine learning," Adsorption Science & Technology, vol. 2022, 2022.

[19] S. Shivaprasad and M. Sadanandam, "Dialect identification using modified features with deep neural networks," Traitement du Signal, vol. 38, no. 6, pp. 1793–1799, Dec. 2021. doi: 10.18280/ts.380622.