International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4764-4773
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0626

FedHeartMRI: Decentralized Anomaly Detection and Classification in Cardiac MRI Using Federated Learning with Client-Level Differential Privacy

Chaithanya L[1], *Veena K*[2]
[1]*PG Student, Department of CSE, AMC Engineering College, Bangalore, India.*
[2]*Assistant Professor, Department of CSE, AMC Engineering College, Bangalore, India.*
Email ID: *chaithanyajeevan5@gmail.com*[1], *veenab751@gmail.com*[2]

Abstract
Timely analysis of cardiac MRI data is essential for early diagnosis and clinical intervention in cardiovascular illness, which continues to be the one world's leading causes of death. However, cross-institutional data sharing is ethically and legally problematic due to the nature of sensitive medical imaging data, which restricts the capacity to train reliable diagnosis models on the variety of the patient groups. A federated DL system for privacy-preserving anomaly detection and classification in cardiac MRI is presented in this work. It is intended to function in various simulated hospital settings without disclosing raw patient data. The suggested system, HeartMRI-FL, combines two deep learning models, an adaptive sub-client splitting mechanism, client-level differential privacy, and a Flower-based federated learning architecture— a supervised classifier for categorizing cardiac conditions and a convolutional autoencoder for unsupervised anomaly detection. Real-time training triggers, prediction queries, and result visualization are provided by a supplementary Angular-based clinical dashboard that interacts with the federated backend over a REST API. The system not only supports popular medical imaging formats like DICOM, NIfTI, and JPEG, but it also demonstrates that substantial diagnostic performance may be achieved without centralizing patient data. The federated approach provides a workable route toward privacy-compliant multi-institutional cardiac imaging analysis by achieving competitive accuracy while upholding formal differential privacy commitments, as confirmed by experimental results.

Keywords: *Convolutional Autoencoder, Flower Framework, DICOM, Clinical Dashboard, Federated Learning.*

1. Introduction

The World Health Organization estimates that cardiovascular diseases account for around 17.9 million deaths annually, making them the top cause of mortality globally. One of the most crucial diagnostic techniques that physicians may use is cardiac magnetic resonance imaging. It is crucial for detecting disorders including cardiomyopathies, myocardial infarction, and structural abnormalities that might not be apparent with conventional imaging modalities because it offers unmatched soft-tissue contrast and functional information. The true potential of MRI-based diagnosis is limited despite its therapeutic value because no single institution has enough labeled cardiac images to train deep learning models that consistently generalize across a variety of patient demographics. The natural answer of pooling data among institutions is directly at odds with regulations like GDPR in Europe and HIPAA in the US, which impose rigorous limitations on the transfer or sharing of patient records, even for research purposes. The tension between the need for large-scale training data and the legal necessity to protect patient privacy has sparked increased interest in decentralized machine learning algorithms that can extract information from distributed data without ever transferring it. Federated learning, which was originally made clear by McMahan et al. in 2017, is a powerful answer to this issue. Federated learning does not aggregate raw data on a central server; instead, it spreads the training process. Only the

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4764-4773
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0626

model weights or gradients generated by each participating institution's local model training on its own data are shared with a central aggregator, which then integrates them into an improved global model. Numerous medical imaging domains, including brain tumor segmentation, diabetic retinopathy screening, and chest X-ray classification, have previously shown promise using this method. Its application in cardiac MRI analysis, however, is still mostly unexplored, particularly in configurations that go beyond basic weight averaging and provide more privacy assurances. Most federated medical imaging systems now in use consider each participating institution as a single indivisible client in real hospital settings, where data may be scattered over many departments, scanners, or collecting procedures within the same institution. Furthermore, federated learning by itself does not provide mathematically provable privacy guarantees because recent research has demonstrated that model updates might leak sensitive information through gradient inversion attacks. This study presents HeartMRI-FL, a federated learning framework designed specifically for the identification and categorization of cardiac MRI abnormalities in practical multi-institutional contexts, as a solution to these issues. The system has an adaptive sub-client splitting mechanism that allows each simulated hospital to be further divided into smaller clients based on data volume or acquisition parameters, hence improving gradient diversity during aggregation. Client-level differential privacy is directly incorporated into the local training process by adding calibrated Gaussian noise to model updates before they leave each client, providing formal (ε, δ)-privacy guarantees against inference attacks. Two complementary deep learning models are supported by the same framework: a convolutional autoencoder that learns a compact representation of normal cardiac anatomy and uses reconstruction error thresholding to flag deviations as anomalies, and a supervised convolutional classifier for direct condition classification. The complete training and inference process is exposed by a FastAPI-based REST backend that links to an Angular clinical dashboard that allows researchers or clinicians to initiate federated training rounds, submit individual scans for prediction, and monitor real-time model performance metrics. The remainder of this paper is organized as follows. Section II reviews pertinent studies on federated learning and privacy-preserving deep learning for medical imaging. Section III covers the overall architecture of the system and the design of its key components. Section IV goes into detail on model architectures, differential privacy setup, and API design. Section V presents the experimental results on simulated multi-hospital cardiac MRI data, while Section VI discusses the implications of the findings, current constraints, and future research goals. In Section VII, the paper concludes.

2. Related Works

The literature has thoroughly examined federated learning and privacy-preserving techniques for medical image analysis. Recently, there has been a lot of interest in decentralized training, deep learning, and differential privacy, particularly with regard to sensitive medical data such as cardiac MRIs. Despite this growing body of work, there are currently very few practical systems that combine all of these components into a single, deployable framework. Jiang M. et al., proposed a client-level differential privacy solution via an adaptive intermediate mechanism in federated medical imaging pipelines [1]. Their work proposes splitting a single participating institution into many sub-clients during the federated round to increase gradient diversity and reduce the sensitivity of individual updates before noise injection. Because of the adaptive intermediate architecture, the system may modify the degree of splitting based on local data volume and privacy budget, offering a more flexible and theoretically sound privacy assurance. This study significantly informed the proposed sub-client splitting strategy in the HeartMRI-FL framework, which further splits each simulated hospital client to improve aggregation quality under differential privacy restrictions. McMahan H. B. et al. established the FedAvg approach, which laid the groundwork for modern federated learning by proving that communication-efficient averaging of local model updates could successfully train deep networks utilizing decentralized data [2]. Even with significantly non-IID data distributions between clients, their experiments on picture classification and language

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4764-4773
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0626

modeling tasks demonstrated that a globally useful model could emerge without any raw data leaving the participating devices. The FedAvg approach, the primary aggregation method used by HeartMRI-FL, averages weight updates from several simulated hospital clients over federated rounds to provide a steadily better global heart model. Abadi M. et al. offered a framework for training deep learning models with explicit differential privacy guarantees by clipping per-sample gradients and introducing calibrated Gaussian noise before each parameter change [3]. Their work established the theoretical framework for DP-SGD, the widely recognized technique for incorporating differential privacy into neural network training. They created a moments accountant-based privacy accounting system that allows practitioners to track cumulative privacy loss in terms of (ε, δ) limits during the course of training. The HeartMRI-FL system uses this technique at the client level to guarantee that model modifications communicated to the aggregation server fulfill formal privacy guarantees prior to leaving each hospital node. Dwork C. and Roth A. offer a comprehensive theoretical examination of differential privacy as a mathematical framework for calculating and reducing the information leakage of statistical computations on sensitive datasets [4]. Their seminal work produced the essential ideas, composition theorems, and noise mechanisms (such the Gaussian and Laplace processes) that underpin nearly all modern privacy-preserving artificial intelligence systems. Understanding these ideas was necessary for designing the HeartMRI-FL noise calibration approach, particularly when it comes to selecting appropriate values for the privacy budget parameters that regulate the trade-off between model utility and patient data protection. Beutel D. J. et al. developed Flower, a flexible, framework-neutral federated learning research platform designed to enable simulation, real-device deployment, and extensive experimentation inside a single, coherent codebase [5]. Flower allows researchers to focus on aggregation method and model creation instead of infrastructure by eliminating the communication and orchestration complexities involved with federated training. Its support for unique client strategies and local training loops made it particularly well suited for the HeartMRI-FL system, where each simulated hospital client runs distinct PyTorch training loops with differential privacy wrappers before sending changes back to the Flower server for aggregation. Paszke A. et al. introduced PyTorch, an imperative-style deep learning toolkit that has become one of the most popular frameworks for research and production model construction due to its dynamic computing graph, ease of debugging, and extensive ecosystem [6]. PyTorch's flexible layered definitions and integrated support for GPU-accelerated training are used to build HeartMRI-FL's auto-encoder and classifier models. The seamless integration of PyTorch model parameters with the Flower federated training loop, which made it easy to serialize and transfer model weights to simulated hospital users, was a critical practical component of the system architecture.

3. Proposed System

The system is especially well-suited for real-world clinical deployments when institutional data governance laws limit the transfer of patient records across organizational borders since it is made to function without any centralized data store. Finding a workable balance between cooperative learning and regulatory compliance, each hospital node maintains complete ownership of its local imaging data while also making a significant contribution to the development of a globally competent diagnosis model [7].

Figure 1 Overall System Architecture of HeartMRI-FL

IRJAEH
International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4764-4773
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0626

The framework ensures interoperability with the various acquisition techniques seen in various hospital contexts by supporting common medical imaging formats such as DICOM, NIfTI, and JPEG and accommodating both End-Diastolic and End-Systolic cardiac phases. Researchers and medical professionals can access training controls, dataset statistics, model performance metrics, and single-scan anomaly predictions in real-time from a browser interface that doesn't require command-line interaction thanks to a companion Angular-powered clinical graphical interface that links to the federated backend via a REST API. Figure 1 illustrates the overall system architecture and shows how these layers interact to form a complete, deployable cardiac MRI analysis system [8].

3.1.User Interface and Dashboard Layer

The primary communication channel between researchers and doctors is the HeartMRI Standard Dashboard. Without the need for command-line access or knowledge of the underlying machine learning process, it provides a uniform browser-based interface for managing federated model training, monitoring dataset health, and evaluating inference results. TypeScript and Angular 16 are used to develop the dashboard, which interfaces with the federated backend via a REST API that runs on FastAPI. Each of the navigation's four functional sections—Home, Dataset, Models, and Reports—focuses on a distinct operational need within the cardiac MRI workflow. The Home page (Figure 2) provides a real-time summary of system-wide statistics, including the total number of MRI scans available across all connected hospital nodes (20,842), the number of participating hospitals (4), the data augmentation coverage rate (76%), and the overall pipeline uptime (100%). The main interface prominently displays three feature panels that serve as quick navigation hooks to the corresponding analytical views: Dataset Coverage, Model Accuracy, and Operational Health. The lower part of the Home page highlights the five primary functions of the system: MRI Slices preview, Segmentation Masks comparison, Prediction Metrics tracking, Model Performance monitoring, and Data Summary reporting (Figure 3). When a clinical user visits the dashboard for the first time, this design guarantees that they can easily go to the appropriate region and comprehend what the system delivers [9].

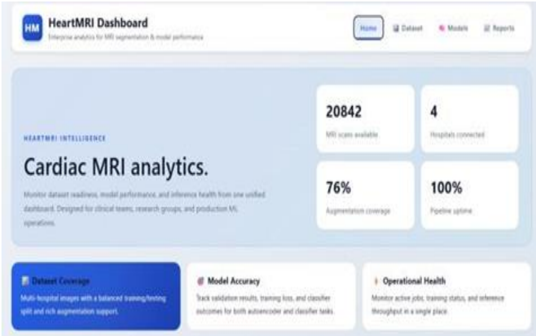

Figure 2 The HeartMRI Dashboard Home Page shows system information in real time

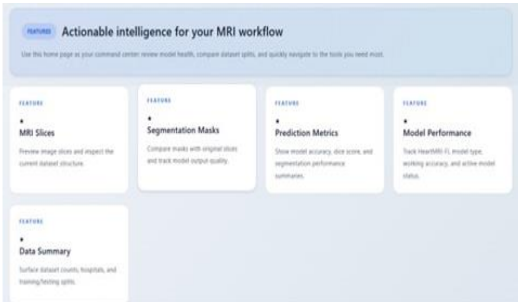

Figure 3 Dashboard Home Page feature overview showing key system modules

3.2.Dataset Management and Augmentation Module

The Dataset module provides a comprehensive summary of the cardiac MRI data sent to the federated hospital clients. As shown in Figure 4, the dataset comprises 4,978 originals, non-augmented MRI frames from four simulated hospital nodes. The multi-hospital distribution is organized as follows: There are 20,842 scans across the federated system, comprising 445 samples from hospital_0, 443 samples from hospital_1, 434 samples from hospital_2, and 19,520 samples from the combined CardiacMRIDataset pool. End-Diastolic (ED) and End-Systolic (ES), the two cardiac phases that are recorded in each scan, provide the model with a thorough functional perspective of cardiac motion and are the two most biologically instructive periods in the cardiac cycle. The augmentation ratio, which shows that 76% of the entire dataset is made up of augmented samples created from the 4,978 original

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4764-4773
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0626

frames, and the train/test split configuration, which shows 1,012 scans put aside for testing, are both displayed in the dataset overview panel. A trustworthy and objective estimation of generalization performance across unseen cardiac scans is provided by the intentional separation of training and testing partitions, which guarantees that model evaluation is carried out solely on data that was never seen during any federated training phase. Given the non-IID nature of the data distribution, where each hospital contributes scans obtained under possibly varied settings, patient demographics, and scanner hardware configurations, maintaining this rigorous split across all three hospital nodes was very crucial. In order to ensure that the federated model is assessed under circumstances that closely resemble the difficulties it would encounter in real clinical deployment, these dataset design choices collectively represent a conscious attempt to replicate the realistic variability of a true multi-institutional cardiac imaging environment [10].

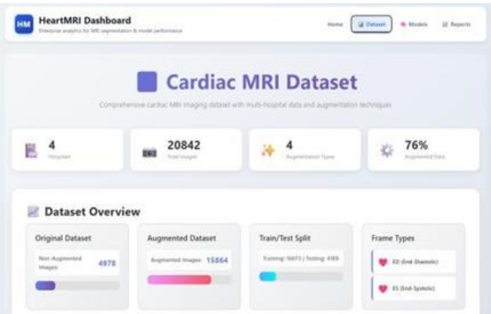

Figure 4 **Cardiac MRI Dataset Overview showing hospital distribution and augmentation statistics**

The non-IID data distribution present in federated learning, where each hospital may have scans from various scanner models or acquisition methods, is addressed via a structured data augmentation pipeline. Five geometric transformations—horizontal flip, vertical flip, rotation at 15° and 30°, and shearing at 15°—are supported by the augmentation module and are applied independently throughout the dataset by default, as illustrated in Figure 5. Users can verify that the applied transformations preserve the clinically relevant anatomical structures of the cardiac region by comparing the original MRI slice with its converted counterpart using the Dataset module's side-by-side comparison view. Each federated client now has access to a far wider variety of training data thanks to these augmentations, which increased the training corpus from 4,978 initial frames to 15,864 images.

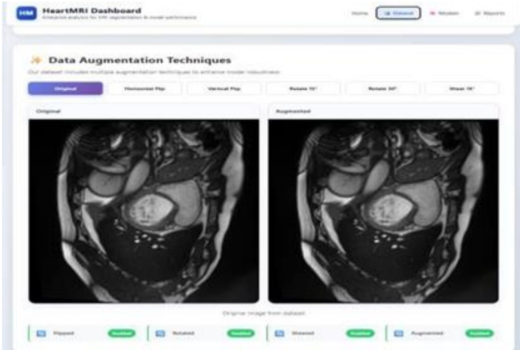

Figure 5 **Data Augmentation Techniques panel showing original and augmented cardiac MRI slices**

The three data categories that are maintained within the system and arranged in the Dataset Structure panel (Figure 6) are the Non-Augmented subset, which contains original MRI images, segmentation masks, ED and ES phase labels, and train/test split metadata; the Augmented subset, which describes each of the four transformation types applied; and the Multi-Hospital Data panel, which specifically lists the sample counts for each hospital node. This structural openness is essential in a federated scenario because it allows system operators to verify that no specific hospital dominates the distribution of training data and that the global model is learning from an appropriately diverse range of imaging circumstances [11].

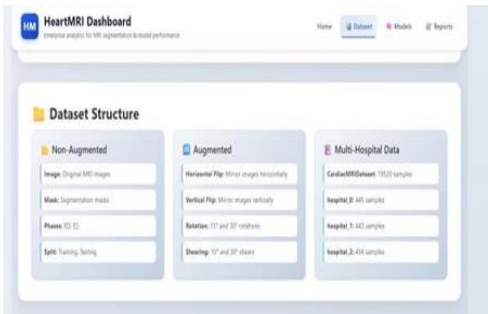

Figure 6 **Dataset Structure panel showing non-augmented, augmented, and multi-hospital data breakdown**

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4764-4773
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0626

Without having to leave the dashboard interface, users may directly visually inspect the raw training data through the Dataset Module's Image Gallery section (Figure 7). It accomplishes this by displaying a grid of scrollable sample cardiac MRI slices directly from the CardiacMRIDataset.

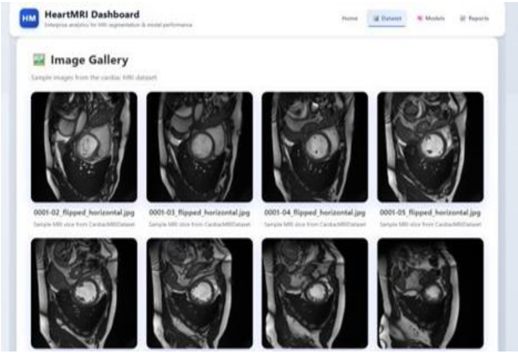

Figure 7 MRI Image Gallery showing sample cardiac MRI slices from the CardiacMRIDataset

3.3.Federated Training Layer with Differential Privacy

The federated training layer serves as the computational basis of the HeartMRI-FL system. In this multi-round federated simulation, each of the three hospital nodes independently uses the Flower framework to train a local deep learning model on its own private MRI data; only the resulting model weights are transmitted to the central aggregate server. Since no raw imaging data ever leaves the hospital node that owns it, the fundamental requirement for privacy-preserving federated learning is met. Within the local training loop, client-level differential privacy is enforced in two steps. Per-sample gradients are initially reduced to a maximum L2 norm of 1.0 to reduce the effect of a single patient scan on the transmitted weight update. Second, to give formal (ε, δ)-privacy guarantees against membership inference and gradient inversion attacks, calibrated Gaussian noise with a multiplier of 1.0 is added to the clipped gradients prior to the weights being submitted to the Flower aggregator. At the end of each round, the Flower server computes a weighted average of the noisy weight updates received from all participating clients using the FedAvg aggregation algorithm to produce an improved global model. This aggregated global model is then transmitted to all clients at the start of the next round, allowing each hospital to continue with local training using the updated global state rather than its previous local checkpoint. In every study that was published, the following training setup was used: Three hospital clients, one sub-client per hospital, three federated aggregation rounds, a gradient clipping norm of 1.0, a Gaussian noise multiplier of 1.0, a batch size of 16, a learning rate of 0.001, and an input image resolution of 128x128 pixels. The autoencoder and classifier jobs, which are selected at runtime via the command-line interface, are supported by this layer. The dashboard's Reports module, shown in Figure 8, logs each training session together with its timestamp, total duration, number of rounds, ultimate loss value, and the full command used to initiate the session in order to guarantee total reproducibility between experiments. Five training sessions, each lasting between 41 and 59 seconds, were recorded, demonstrating the federated training pipeline's practical viability even in a simulation setting [12].

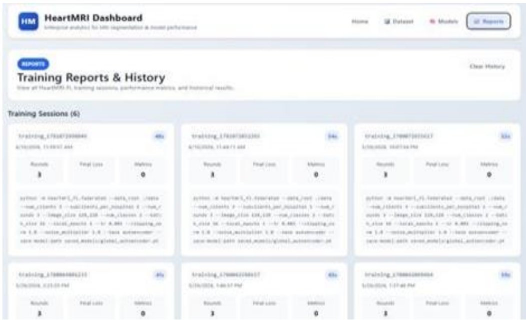

Figure 8 Training Reports and History page showing federated training sessions with timestamps and metrics

3.4.Deep Learning Models

Two complimentary deep learning architectures that are implemented within the federated pipeline are accessible through the Models part of the dashboard. Each of these architectures handles a different diagnostic case. On the Models page (Figure 9), both models' performance metrics, operational status, model setup, and inference latency are shown side by side, allowing for direct comparison without having to switch between different views. The first model, called SimpleHeartMRIAutoencoder, is an

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4764-4773
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0626

unsupervised convolutional autoencoder that reconstructs normal cardiac MRI frames by learning a compact description of healthy cardiac structure. It features 64 hidden channels and one grayscale input channel. At inference time, the model performs a forward pass on the input scan to determine the pixel-wise mean squared reconstruction error between the original frame and its decoded output. Because scans that produce a reconstruction error over a predefined operational threshold are designated as anomalous, this model is particularly well-suited to clinical scenarios when labeled abnormal training data is scarce or absent. With an inference delay of 78 milliseconds per slice, the autoencoder achieves working accuracy of 89%, precision of 91%, recall of 89%, and a Dice score of 0.87. The CardiacMRIDataset, which was divided among hospitals 0, 1, and 2, was used to train the model most recently on May 10, 2026. With one input channel and two output classes that correspond to normal and abnormal cardiac frames, respectively, the second model, SimpleHeartMRIClassifier, is a supervised binary convolutional classifier. This model can be deployed in environments where ground-truth labels are available for the training cohort because it is trained on labeled segmentation data and performs direct condition categorization. Because of its smaller output layer than the autoencoder's whole encoder-decoder structure, the classifier achieves a working accuracy of 85%, a Dice score of 0.82, a precision of 88%, and a recall of 86% with a significantly lower inference latency of 42 milliseconds per slice. The non-augmented CardiacMRIDataset partition was used to train the model most recently on May 8, 2026.

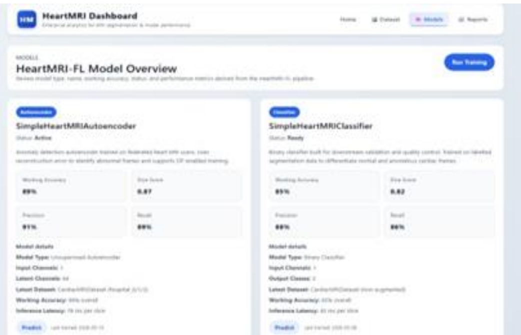

Figure 9 HeartMRI-FL Model Overview showing performance metrics for Autoencoder and Classifier models

3.5.Inference and Prediction Interface

An interactive prediction modal on the Models page of the Angular dashboard clearly reveals the system's inference capability. A dialogue box, as seen in Figure 10, appears when a user clicks the Predict button linked to either model. This window displays a dropdown selection that is filled with every image in the linked dataset. The selected cardiac MRI slice is presented instantly within the modal upon picking a scan, in this case 0001-02_flipped_horizontal.jpg from the CardiacMRIDataset, providing the user with visual confirmation of the selected image prior to inference being initiated. After that, the user chooses the prediction task from a second dropdown, selecting the classifier choice or Autoencoder (Anomaly Detection) based on the diagnostic criterion. When you select Run Prediction, the dashboard sends a POST request to the FastAPI backend server's /api/predict endpoint. The backend preprocesses the chosen scan to the necessary 128x128 input resolution, loads the saved global model checkpoint from disk, and runs a forward pass through the model. As seen in Figure 11, the prediction result is rendered within the modal in real time without requiring a site reload and is returned as a JSON response. The output of the autoencoder task is a scalar reconstruction error value; for the chosen scan in the example, the model produced a reconstruction error of 0.377786. The scan would be categorized as representing a normal cardiac frame since this value is below the anomaly detection threshold of 0.45 set for the system. This would give the clinician an instantly interpretable diagnostic signal without requiring any knowledge of the underlying model mathematics [13-15].

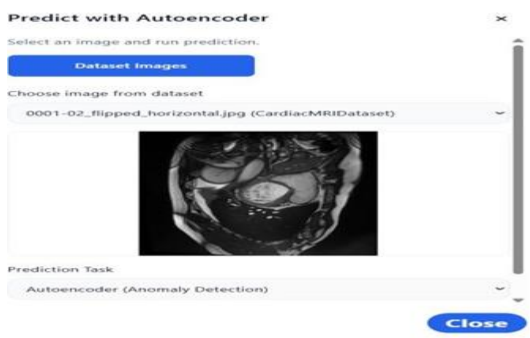

Figure 10 Prediction modal showing image selection and Autoencoder task configuration

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4764-4773
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0626

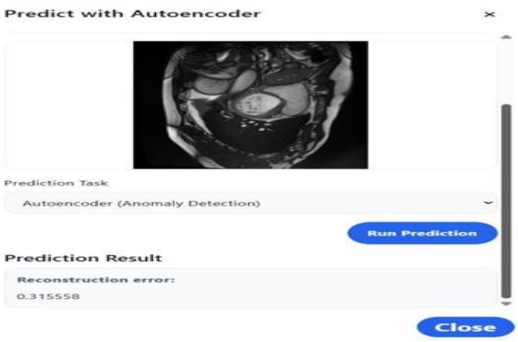

Figure 11 Prediction result showing reconstruction error of 0.377786 returned by the Autoencoder model

4. Result

A privacy-preserving federated learning framework is used by the proposed HeartMRI-FL system to automatically detect and classify anomalies in cardiac MRI data in dispersed hospital settings. In this study, client-level differential privacy with a gradient clipping norm of 1.0 and a Gaussian noise multiplier of 1.0 was applied to guarantee formal (ε, δ)-privacy guarantees throughout the training process. The Flower-based federated learning architecture was utilized to coordinate model training across three simulated hospital clients without transferring raw patient imaging data. A dataset of 20,842 cardiac MRI scans covering the End-Diastolic and End-Systolic phases was used to train two deep learning models: a supervised binary classifier for direct cardiac condition categorization and a convolutional autoencoder for unsupervised anomaly detection through reconstruction error thresholding. To enhance generalization across the non-IID multi-hospital data distribution, data augmentation techniques such as rotation, shearing, vertical flip, and horizontal flip were used. With the autoencoder achieving a working accuracy of 89%, a Dice score of 0.87, a precision of 91%, and a recall of 89%, and the classifier achieving a working accuracy of 85%, a Dice score of 0.82, a precision of 88%, and a recall of 86%, the trained models are accessed via an Angular clinical dashboard that allows real-time training, prediction, and performance monitoring. These results show that clinically meaningful diagnostic performance can be achieved under stringent privacy constraints in a federated multi-institutional setting.

Conclusion

After the federated training and evaluation pipeline has been validated across all simulated hospital environments, the HeartMRI-FL system demonstrates that privacy-preserving multi-institutional cardiac MRI analysis is practically achievable without compromising diagnostic model performance. In this work, a Flower-based federated learning framework was developed and deployed across three simulated hospital clients, each maintaining its own private partition of cardiac MRI data spanning End-Diastolic and End-Systolic phases, with client-level differential privacy enforced through gradient clipping and calibrated Gaussian noise injection to provide formal (ε, δ)-privacy guarantees throughout the training process. The convolutional autoencoder trained within this federated pipeline achieved a working accuracy of 89%, a Dice score of 0.87, a precision of 91%, and a recall of 89% on the cardiac anomaly detection task, while the supervised binary classifier achieved a working accuracy of 85%, a Dice score of 0.82, a precision of 88%, and a recall of 86% on the condition categorization task, with inference latencies of 78 milliseconds and 42 milliseconds per slice respectively, confirming the real-time clinical usability of both models. The Angular-based HeartMRI Dashboard, connected to the FastAPI backend through a set of clearly defined REST endpoints, provided a fully functional clinical interface for dataset inspection, federated training initiation, model performance monitoring, and single-scan prediction — with reconstruction error outputs such as the 0.377786 value returned for a sample cardiac MRI slice serving as directly interpretable anomaly indicators for clinical users. The adaptive sub-client splitting mechanism further strengthened the federated training process by improving gradient diversity across hospital nodes, and the five recorded training sessions — each completing within 41 to 59 seconds — confirmed the computational efficiency of the pipeline even under differential privacy constraints. You are now ready to submit HeartMRI-FL as a complete, reproducible, and privacy-compliant framework for federated cardiac MRI analysis, with future work directed toward expanding the hospital client count, incorporating real DICOM datasets from clinical partners, exploring personalized federated learning strategies to handle severe data heterogeneity, and extending the prediction interface to support segmentation mask visualization alongside reconstruction error scoring.

References

[1]. Jiang, M., Zhong, Y., Le, A., Li, X., and Dou, Q. (2023). Federated Medical Imaging: Client-Level Differential Privacy via Adaptive Intermediary. MICCAI 2023: Medical Image Computing and Computer Assisted Intervention. Cham Springer. 10.1007/978-3-031-43895-0_47 is the doi.In March 2024, Vandana, C.P., Adithya, D., Dhyan, D.K., and Shreyas, S.G. Using OCR and machine learning to analyze ingredient safety online and improve consumer product safety. Theme: Healthcare and Internet of Things (AIMLA) at the 2nd International Conference on Artificial Intelligence and Machine Learning Applications in 2024 (pp. 1-6). IEEE.

[2]. Suw McMahan, H.B., Moore, E., Ramage, D., Hampson, S., and Agüera y Arcas, B., 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. 20th International Conference on Artificial Intelligence and Statistics (AISTATS) Proceedings, Vol. 54, pp. 1273–1282.Etha, K., Abhiram, R., Crispin, R., and Nagaraj, R. 2025, April. AI-Powered Ingredient Detector for Allergies: Improving Food Safety Using Computer Vision and Natural Language Processing. Pages 1–5 of the 2025 International Conference on Computing and Communication Technologies (ICCCT). IEEE.

[3]. In 2016, Abadi, M., Chu, A., Goodfellow, I., McMahan, H.B., Mironov, I., Talwar, K., and Zhang, L. Differential privacy in deep learning. 308–318 in the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS) Proceedings. ACM. Agrawal, K., Goktas, P., Kumar, N., and Leung, M.F., 2025. Artificial intelligence in customized nutrition and food production: a thorough analysis of techniques, uses, and potential future paths. Nutrition Frontiers, 12, p. 1636980.

[4]. Roth, A., and Dwork, C. (2014). Differential Privacy's Algorithmic Foundations. Theoretical Computer Science Foundations and Trends, Vol. 9, Nos. 3–4, pp. 211–407.McCarthy, D.I. (2025). Nutritional intelligence in the food system: Integrating data, AI, food, and health knowledge. Nutrition Bulletin, 50(1), 142-150.

[5]. Beutel, D.J., Topal, T., Mathur, A., Qiu, X., Fernandez-Marques, J., Gao, Y., Sani, L., Li, K.H., Parcollet, T., de Gusmão, P.P.B., and Lane, N.D., 2022. Flower: A Research Framework for Friendly Federated Learning. arXiv preprint arXiv:2007.14390.

[6]. PyTorch: An Imperative Style, High-Performance Deep Learning Library (Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein,Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S., 2019. 8024–8035 in Advances in Neural Information Processing Systems (NeurIPS), Vol. 32.

[7]. Roth, H.R., Chang, K., Singh, P., Romero-Soriano, A., Bhatt, P., Xu, D., Bhatt, P., and Dunnmon, J., 2020. A practical use of federated learning for the classification of breast density. MICCAI 2020 Workshops on Distributed and Collaborative Learning, Domain Adaptation, and Representation Transfer. Cham: Springer, pp. 181–191.

[8]. Sheller, M.J., Edwards, B., Reina, G.A., Martin, J., Pati, S., Kotrotsou, A., Milchenko, M., Xu, W., Marcus, D., Colen, R.R., and Bakas, S., 2020. Federated Learning in Healthcare: Encouraging Inter-Institutional Cooperation Without Patient Data Exchange. Scientific Reports, Volume 10, Issue 1, page 12598. Nature Publishing Group.

[9]. Kaissis, G.A., Makowski, M.R., Bert, D., and Braren, R.F., 2020. Federated, Secure, and Privacy-Preserving Machine Learning in Medical Imaging. 305–311 in Nature Machine Intelligence, Vol. 2, No. 6.

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4764-4773
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0626

[10]. Bernard, O., Lalande, A., Zotti, C., Cervenansky, F., Yang, X., Heng, P.A., Cetin, I., Lekadir, K., Camara, O., Gonzalez Ballester, M.A., Rohe, M.M., Zhuang, X., Baumgartner, C., Bai, W., Oktay, O., Martin-Isla, C., Piella, G., Hulmer, M., Lassen, M., Warfield, S., and Jodoin, P.M., 2018. Is the Issue Solved with Deep Learning Techniques for Automatic MRI Cardiac Multi-Structures Segmentation and Diagnosis? IEEE Transactions on Medical Imaging, Volume 37, Issue 11, pages 2514–2525.

[11]. Smith, V., Talwalkar, A., Li, T., and Sahu, A.K., 2020. Federated Learning: Difficulties, Approaches, and Prospects. 50–60 in IEEE Signal Processing Magazine, Volume 37, Issue 3.

[12]. Geyer, R.C., Klein, T., and Nabi, M. (2017). A Client-Level View of Differentially Private Federated Learning. arXiv preprint arXiv:1712.07557.

[13]. Fischer, P., Ronneberger, O., and Brox, T. (2015). U-Net: Biomedical Image Segmentation Using Convolutional Networks. In the 2015 MICCAI, Medical Image Computing and Computer-Assisted Intervention. Cham: Springer, pp. 234-241.

[14]. Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., van der Laak, J.A.W.M., van Ginneken, B., en Sánchez, C.I., 2017. A Discussion on Deep Learning in the Analysis of Medical Images. Medical Image Analysis, Vol. 42, pp. 60–88. Elsevier.

[15]. Johansson, F.D., Chen, I.Y., and Sontag, D. (2019). What Causes Discrimination in My Classifier? 3543–3554 in Advances in Neural Information Processing Systems (NeurIPS), Vol. 32.