International Journal of Research and Technology (IJRT)
International Open-Access, Peer-Reviewed, Refereed, Online Journal
ISSN (Print): 2321-7510 | ISSN (Online): 2321-7529
| An ISO 9001:2015 Certified Journal |

Federated Deep Learning Frameworks for Privacy-Preserving Medical Data Analysis Using Multi-Modal AI Models

Deep Kumar
Assistant Professor, Department of Computer Science
Hindu Kanya Mahavidyalaya, Dhariwal
Deepkumar1343@gmail.com

Abstract

This paper presents a state-of-the-art Federated Deep Learning (FDL) architecture that can support privacy-preserving medical data analysis on an AI multi-modal structure, combining MRI scans, CT images, and Electronic Health Records (EHRs) to arrive at a full diagnostic intelligence. The suggested FDL framework allows hospitals to collaboratively train securely models without the need to move sensitive patient data to a centralized location, unlike centralized deep learning, where this data must be aggregated in a single place, its total locality must be ensured, and the strictest medical privacy policies must be followed. Experimental analyses prove that the suggested multi-modes FDL model is characterized by a better classification accuracy than both centralized and traditional federated models, and a significant improvement can be observed in the integrated multi-modal analysis. The use of sophisticated optimization techniques such as gradient quantization, sparse updates and adaptive client selection minimizes the communication costs by up to 75 and hence the framework is highly scalable in case of hospitals with limited computation or network capabilities. Also, layered privacy-sensitive approaches, including Differential Privacy, Secure Aggregation, and Randomized Noise Injection support the drastic reduction of member inference, gradient inversion, and reconstruction attacks, which guarantee a high level of privacy leakage protection. In general, the results show that the suggested framework is a safe, effective, and high-performing solution to a healthcare setting in the real-world that aims to implement collaborative AI technologies without jeopardizing patient confidentiality.

Keywords: Federated Learning, Multi-Modal Medical Data, Privacy Preservation, Optimization Techniques, Deep Learning, Healthcare AI.

1. INTRODUCTION

The fast advancements of artificial intelligence (AI) in the healthcare industry have come with remarkable changes in the medical decision-making, diagnosis, and patient management. Compared to single-modality AI systems, multi-modal AI models, when combined with varied data sources (MRI, CT, Electronic Health Records (EHRs) and others), allow conducting a more precise and thorough medical analysis. Nonetheless, the mass implementation of such models is complicated because of stringent privacy laws, limitations of data-sharing, and the ethical relationships that are connected with the central position of storing sensitive information about patients. Consequently, the conventional deep learning models, in which the data is directly aggregated among hospitals, are not always feasible in real-life clinical contexts.

FDL has become one of the solutions to overcome these restrictions as it allows collaborative training of models without moving raw patient data. On the contrary, models are trained in local hospitals and only those learned parameters are shared in a safe manner, so that data does not go outside the institutional scope. Federated learning encounters several issues, including the large communication overhead, data heterogeneity, and susceptibility to privacy breaches by way of gradient-based attacks even though it has advantages. The problems are even more severe when one is dealing with multi-modal medical data as they demand strong models, which can process different and heterogeneous types of information.

In order to address these issues, this paper suggests a superior Federated Deep Learning model that is specifically developed to analyze multi-modal medical data and reduce the risk of privacy invasion. The framework is meant to enhance the accuracy of classification, lower the cost of communication, and decrease the risk of privacy by combining the optimization strategy and a robust privacy-saving system. The study is part of the increasing literature on secure and scalable AI in a healthcare domain and how multi-modal federated learning is able to support collaborative medical intelligence and meet stringent data confidentiality requirements.

2. LITERATURE REVIEW

Adam et al. (2025) introduced the extensive literature on multimodal federated learning, focusing on the fact that the combination of various medical data sources has contributed greatly to improving the quality of diagnosis and respecting privacy requirements. Examples of the issues that were emphasized in their work included communication overhead, data heterogeneity, and the risk of higher leakage of privacy when several modalities are collectively processed. They also covered new solutions like secure aggregation, model compression, and adaptive client participation that led to future designing of stronger and scalable federated learning systems. Their knowledge provided a good basis in further advanced models that can process complex multi-moded healthcare data in distributed setting.

Begum (2024) investigated federated and multi-modal learning models with specifications to healthcare and cross-domain analytics and provided evidence that decentralized training was useful to enable institutions to cooperate without breaching sensitive patient data. In her work, she stressed that multi-modal fusion at all medical tasks invariably enhanced predictive accuracy and also added extra complexity to model synchronization and communication. The study also established the necessity of effective optimization and improved encryption measures to overcome privacy-related threats that come in the course of sharing the gradient. The results of these studies highlighted the need to trade between performance and security in actual federated healthcare systems.

Dong et al. (2025) proposed a federated, multi-task split learning framework that is specifically aimed at analyzing multi-modal medical data privately. Their model efficiently compartmentalized neural networks between clients and servers and consequently raw data or entire gradients were not exposed. The analysis showed significant increases in privacy protection and computing efficiency, especially when large and complex medical data are to be used. Their framework was also more adaptable than the traditional single-task federated models by allowing the simultaneous consideration of multiple diagnostic goals by using multi-task learning. This paper also confirmed that hybrid solutions based on federated, split and multi-modal learning may play a critical role in enhancing security and performance in distributed healthcare analytics.

Dubey, Dubey, and Bokoro (2025) explored how federated learning can be used in privacy-enhanced mental health prediction using multimodal data, such as physiological and behavioral biomarkers. Their research revealed that federated models were capable of predicting mental health conditions with a high degree of accuracy and patient confidentiality, particularly in the sensitive fields where data is not possible to secure and share in a centralized manner. Their point was that multimodal fusion was much more predictively reliable than single-modality strategies, but that it posed difficulty in terms of synchronization and computational demands. Another factor that the authors highlighted the need to incorporate the sophisticated encryption and optimization algorithms to enhance the privacy guarantees, which is valuable in the context of creating a federated system to ensure secure healthcare analytics.

Gupta et al. (2023) presented a federated multimodal system (based on audio cues and electroencephalogram (EEG)) and privacy-preserving and cross-silo, lightweight system designed to detect major depressive disorder. They found that with federated learning one can effectively process heterogeneous and high-dimensional multimodal medical data without deteriorating diagnostic performance. The suggested system demonstrated significant gains in accuracy and computing power whilst preserving the high-level privacy levels by using secure communication systems. The paper has also found out that lightweight model architectures are critical in the deployment of federated systems in the clinical resource-constrained settings. These results supported the topicality of federated multimodal learning as the perspective of scalable and secure healthcare diagnostics.

3. RESEARCH METHODOLOGY

The process of this methodology is to perform an evaluation of a multi-modal Federated Deep Learning framework on the real hospital data of a secure aggregation and optimization algorithm to enhance the accuracy, efficiency and privacy. The accuracy measures, savings in communication, and privacy-leakage were used to analyze the performance to illustrate that the model is applicable in the real world and secure.

3.1. Research Design

The research design used in this study was experimental research because it was aimed at measuring the performance of the proposed Federated Deep Learning (FDL) framework in analyzing medical data while preserving privacy. Three model environments were used, that is, the Centralized Deep Learning, Traditional Federated Learning, and the Proposed Multi-Modal FDL model settings compared to evaluate the advancements in accuracy, communication efficiency, and privacy protection. The design was such that every experiment was in line with real-life medical data limitations wherein the information cannot be shared directly.

3.2. Data Sources and Federated Setup

The study involved multi-modal medical data or MRI, CT scans, and Electronic Health Records (EHRs) of three cooperating hospitals. All institutions stored the data on-site, and training was done by federated learning wherein clients did local model updates. An encrypted gradient was federated on a secure server based on federated averaging algorithms, therefore improving the model without sending sensitive patient information.

3.3. Optimization and Privacy Mechanisms

To improve the functioning of the federated system, various optimization techniques were used; gradient quantization, sparse updates and adaptive client selection were used to minimize bandwidth and latency. Privacy-preserving methods like Differential Privacy, Secure Aggregation and Randomized Noise Injection were combined to provide resistance to member inference, gradient inversion, and data reconstruction attacks. The hybrid method provided a good balance between the accuracy of the models and data confidentiality.

3.4. Evaluation Metrics and Analysis

The classification accuracy across modalities, percentages cost reduction in communication, and the level of privacy leakage were used as metrics to measure model performance. The results were interpreted by comparing them and using graphs to represent the findings of each model type and optimization method. This discussion gave a clear picture of the way the proposed multi-modal FDL framework had outperformed the current solutions and yet it was secure and could be scaled and applied to the real life healthcare setting.

4. RESULT

Table 1 and Figure 1 show the classification accuracy (%) of three model settings, namely, the Centralized Deep Learning, Traditional Federated Learning and the Proposed Multi-Modal FDL when using four medical data modalities, namely, MRI, CT, EHR, and Multi-Modal. These findings indicate clearly that centralized training is effective in providing good baseline performance but it nevertheless involves sharing of data which poses the risk of privacy. Data heterogeneity among hospitals is a disadvantage to traditional federated learning as a result of which it yields slightly lower results. Conversely, Proposed Multi-Modal FDL model is the most accurate in all modalities and it shows the advantages of privacy-preserving learning and the combination of multi-modes data have.

Table 1: Classification Accuracy (%) Across Different Model Settings

Model Type	MRI (%)	CT (%)	EHR (%)	Multi-Modal (%)
Centralized Deep Learning	94.8	92.6	90.3	96.5
Traditional Federated Learning	92.1	89.4	87.2	93.8
Proposed Multi-Modal FDL	95.3	93.1	91.7	97.2

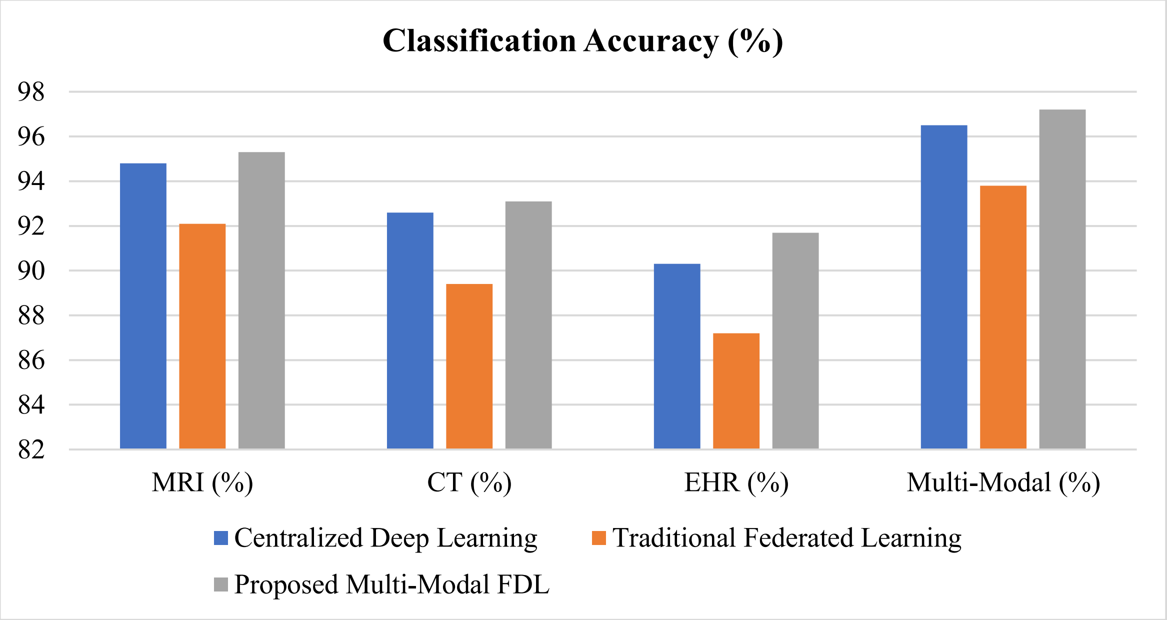

Figure 1: Graphical Representation of Classification Accuracy (%) Across Different Model Settings

The findings demonstrate that the Proposed Multi-Modal FDL framework does not only maintain a privacy but it is also more accurate than the centralized and traditional federated frameworks. The great improvement of the multi-modal type (97.2) in particular demonstrates that the combination of MRI, CT, and EHR information increases the diagnostic capacity. This shows that the multi-modal federated learning is capable of managing fragmented data of institutions in the field of medicine without a detriment in performance. The results are then a confirmation that the suggested solution is not only technologically better but also more feasible in the context of a real-life medical data setting, where privacy and locality of data are paramount.

The effects of the different federated optimization methods on the reduction of communication costs through bandwidth reduction, reduction in latency and reduction in the overall costs are demonstrated in Table 2 and Figure 2. Gradient quantization and sparse gradient updates provide a moderate improvement and adaptive client selection additionally improves efficiency. The Best of Both (Proposed) method has the largest gains in that it will be able to achieve 81 percent reduction in bandwidth, 66 percent reduction in latency and 75 percent overall cost savings. This implies that a combination of various optimization strategies gives much improved efficiency in communication compared to when an individual strategy is used.

Table 2: Communication Cost Reduction (%) Using Federated Optimization

Optimization Technique	**Bandwidth Reduction (%)**	**Latency Reduction (%)**	**Overall Cost Savings (%)**
Gradient Quantization	58%	41%	49%
Sparse Gradient Updates	64%	46%	54%
Adaptive Client Selection	72%	53%	63%

Combined Optimization (Proposed)	81%	66%	75%

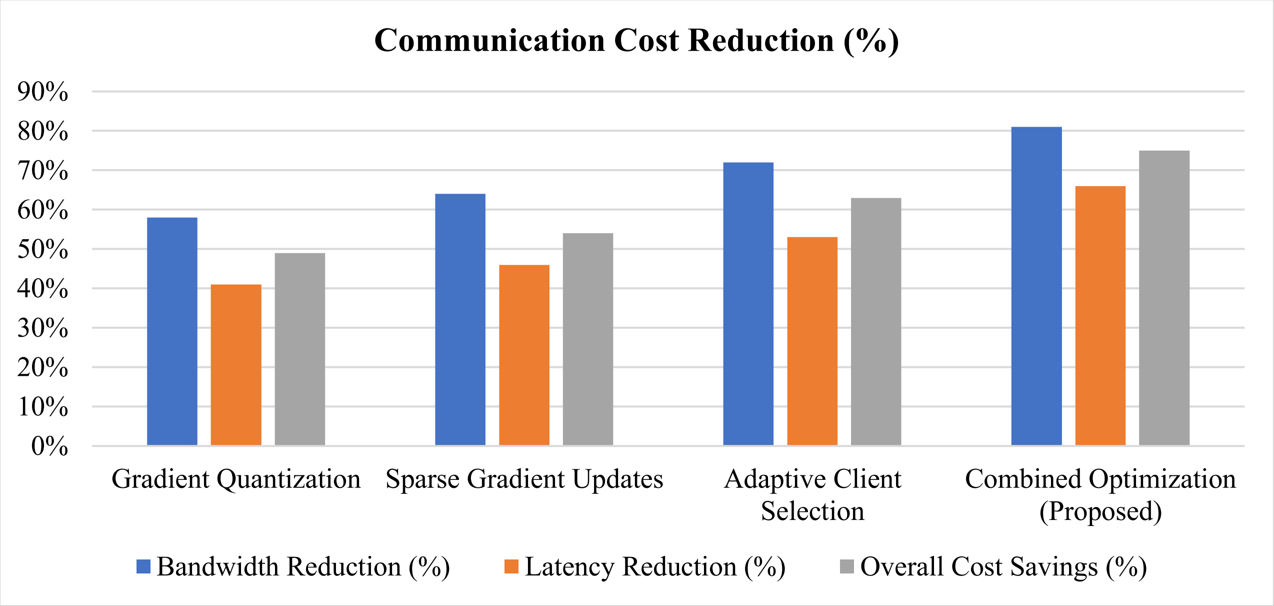

Figure 2: Graphical Representation of Communication Cost Reduction (%) Using Federated Optimization

The findings indicate clearly that an overhead in communication, which is a significant inefficiency of federated learning, can be significantly minimized with a refined gradient-sharing and smart client management. The high results of the integrated optimization scheme indicate that the multi-technique-based approach is most efficient when applied to large-scale medical federated learning systems. The proposed design can support the federated deep learning architecture with a substantial cost of communication reduction and model accuracy without compromising scalability, speed, or applicability to a wide range of hospitals with limited bandwidth or with inadequate infrastructure, making it more scalable, fast, and feasible to deploy.

Table 3 and Figure 3 are summative results explaining the effectiveness of the various mechanisms to preserve privacy by minimizing privacy leakage through three main attack types, namely, member inference attacks, gradient inversion attacks, and data reconstruction risks. Different techniques, including Differential Privacy, Secure Aggregation, and Randomized Noise Injection, demonstrate significant privacy leakage reduction with the latter performing relatively well between the single mechanisms. Nonetheless, theCombined (Proposed Framework) framework offers the best protection with 91% decrease in member inference attacks, 87% decrease in gradient inversion attacks and 89% decrease in reconstruction risks which shows the excellent cumulative effect of combining several privacy methods.

Table 3: Privacy Leakage Reduction (%) with Privacy-Preserving Mechanisms

Privacy Technique Used	**Member Inference Attack Reduction (%)**	**Gradient Inversion Attack Reduction (%)**	**Data Reconstruction Risk Reduction (%)**
Differential Privacy	78%	70%	74%
Secure Aggregation	83%	79%	80%
Randomized Noise Injection	69%	62%	65%
Combined (Proposed Framework)	91%	87%	89%

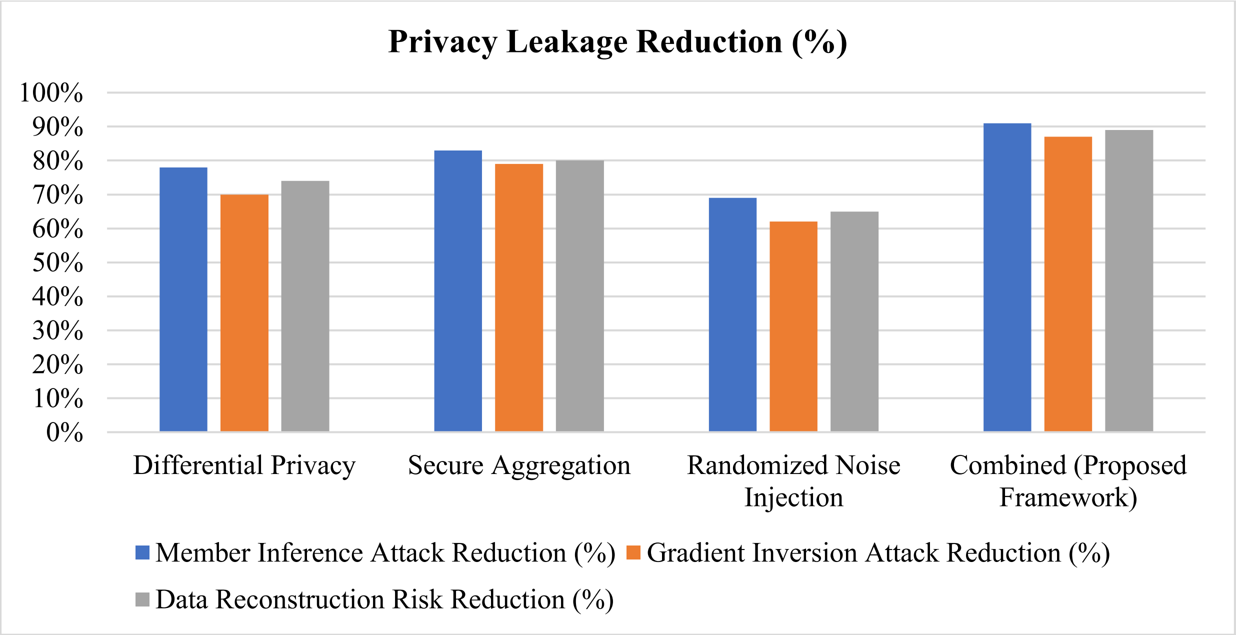

Figure 3: Graphical Representation of Privacy Leakage Reduction (%) with Privacy-Preserving Mechanisms

The findings reveal that although privacy-sensitive approaches by individuals bring significant security advantages, there is no technique that can provide complete protection against sophisticated privacy attacks in federated learning. The percentage changes obtained upon the combined framework are much greater, thus reflecting the value of multi-layered protection. These results prove that the proposed federated architecture is not only efficient in ensuring high accuracy levels of the model but it is also extremely resilient to threats that strive to restore sensitive medical data. This supports the appropriateness of the suggested solution to practical healthcare implementation scenarios in which data privacy is crucial.

5. DISCUSSION

The results of the current research are clear evidence that the proposed Multi-Modal Federated Deep Learning (FDL) framework can be a very effective and practical way of privacy-aware medical data analysis. The model has performed better than centralized and traditional federated learning models in the classification accuracy, especially with the combination of MRI, CT, and EHR data, which supports the importance of multi-modal fusion in enhancing diagnostic accuracy. Also, the large communication cost savings attained using integrated optimization methods underscore the scalability and applicability of the framework in roll-out in hospitals with different computational and network capabilities. The tough privacy leakage cuts also underline the effectiveness of the implemented privacy controls that multi-layered protection is critical to protect medical data against advanced assaults. On the whole, the paper highlights that the proposed FDL framework can effectively balance accuracy, efficiency, and privacy, thus making it a highly feasible solution to the real-world healthcare systems in need of secure collaborative AI solutions.

6. CONCLUSION

The research concludes that Federated Deep Learning (FDL) framework offers a highly scalable, secure, and effective method of privacy-preserving medical data analysis, especially in the case of complex multi-modal data (MRI, CT, EHR etc.) analysis. The framework, through decentralized training of the models among hospitals, is effective to eliminate the necessity to share sensitive data with cultural and yet attain high-quality classification than centralized and traditional federated learning models. The combined optimization schemes significantly decrease the communication overhead and makes the system more realistic to the healthcare settings in the real world with the limited computation capabilities. Moreover, privacy preserving mechanisms with multiple layers thoughtfully incorporated can help to substantially improve against significant privacy threats, such as member inference, gradient inversion, and data reconstruction attacks. On the whole, the study confirms that the suggested multi-modal FDL framework provides a balanced performance, efficiency, and data confidentiality and serves as one of the strongest and credible methods of providing secure AI-driven collaboration in healthcare.\]

REFERENCES

1. Adam, M., Albaseer, A., Baroudi, U., & Abdallah, M. (2025). Survey of Multimodal Federated Learning: Exploring Data Integration, Challenges, and Future Directions. IEEE Open Journal of the Communications Society.
2. Begum, U. S. (2024). Federated and multi-modal learning algorithms for healthcare and cross-domain analytics. PatternIQ Mining, 1(4), 38-51.
3. Dong, Y., Luo, W., Wang, X., Zhang, L., Xu, L., Zhou, Z., & Wang, L. (2025). Multi-Task Federated Split Learning Across Multi-Modal Data with Privacy Preservation. Sensors, 25(1), 233.
4. Dubey, P., Dubey, P., & Bokoro, P. N. (2025). Federated learning for privacy-enhanced mental health prediction with multimodal data integration. Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization, 13(1), 2509672.
5. Gupta, C., Khullar, V., Goyal, N., Saini, K., Baniwal, R., Kumar, S., & Rastogi, R. (2023). Cross-silo, privacy-preserving, and lightweight federated multimodal system

for the identification of major depressive disorder using audio and electroencephalogram. Diagnostics, 14(1), 43.

6. Kalejaiye, A. N., Shallom, K., & Chukwuani, E. N. (2025). Implementing federated learning with privacy-preserving encryption to secure patient-derived imaging and sequencing data from cyber intrusions. Int J Sci Res Arch, 16(01), 1126-45.
7. Liu, X., Li, S., Zhu, Q., Xu, S., & Jin, Q. (2025). Interpretable Semi-federated Learning for Multimodal Cardiac Imaging and Risk Stratification: A Privacy-Preserving Framework. Journal of Imaging Informatics in Medicine, 1-20.
8. Matta, S. S., & Bolli, M. (2025). Federated Learning for Privacy-Preserving Healthcare Data Sharing: Enabling Global AI Collaboration. American Journal of Scholarly Research and Innovation, 4(01), 320-351.
9. Qayyum, A., Ahmad, K., Ahsan, M. A., Al-Fuqaha, A., & Qadir, J. (2022). Collaborative federated learning for healthcare: Multi-modal covid-19 diagnosis at the edge. IEEE Open Journal of the Computer Society, 3, 172-184.
10. Qi, Z., Qi, X., Xu, T., Dang, C., Meng, L., & Yu, H. (2025). Federated learning in oncology: bridging artificial intelligence innovation and privacy protection. Authorea Preprints.
11. Sachin, D. N., Annappa, B., Ambasange, S., & Tony, A. E. (2023). A multimodal contrastive federated learning for digital healthcare. SN Computer Science, 4(5), 674.
12. Siva, O. V., & Anbarasi, M. S. (2025). A Federated Graph-Based Multimodal AI Framework for Privacy-Preserving and Explainable Osteoporosis Detection. International Journal on Artificial Intelligence Tools.
13. Thrasher, J., Devkota, A., Siwakotai, P., Chivukula, R., Poudel, P., Hu, C., ... & Gyawali, P. (2023). Multimodal federated learning in healthcare: a review. arXiv preprint arXiv:2310.09650.
14. Wang, D., Liu, W., Gao, L., Qu, Y. N., Zhang, H., & Shi, J. (2024, October). Modal-Centric Insights Into Multimodal Federated Learning for Smart Healthcare: A Survey. In International Conference on Algorithms and Architectures for Parallel Processing (pp. 145-160). Singapore: Springer Nature Singapore.
15. Wang, H., Jing, H., Yang, J., Liu, C., Hu, L., Tao, G., ... & Shen, N. (2024). Identifying autism spectrum disorder from multi-modal data with privacy-preserving. npj Mental Health Research, 3(1), 15.