Journal of Neonatal Surgery
ISSN(Online): 2226-0439
Vol. 14, Issue 5 (2025)
https://www.jneonatalsurg.com
OPEN ACCESS

Federated Learning Enabled Wireless Sensor Architecture for Secure and Intelligent Brain Surgery Monitoring

Dr. R. Nallakumar[1], Dr. T. Parameswaran[2], Dr. P. M. Benson Mansingh[3], Dr. A. Britto Manoj[4], Dr. S. Lavanya[5]

[1]Associate Professor, Department of Artificial Intelligence and Data Science, Karpagam Institute of Technology, Coimbatore – 641105.

Email ID: drnallakumar@gmail.com

[2]Associate Professor, Department of Computer Science and Engineering, School of Engineering and Technology, CMR University, Bagalur, Bangalore – 562 149

Email ID: parameswaranthangaraj@gmail.com

[3]Assistant Professor (Senior Level), Department of Advanced Computer Science and Engineering, Vignan's Foundation for Science, Technology & Research, Guntur,Andhra Pradesh.

Email ID: drbm_acse@vignan.ac.in

[4]Assistant Professor, Department of Advanced Computer Science and Engineering, Vignan's Foundation for Science, Technology & Research, Guntur,Andhra Pradesh.

Email ID: drbr_acse@vignan.ac.in

[5]Associate Professor, Department of Computer Science and Engineering, Sri Ranganathar Institute of Engineering and Technology, Athipalayam, Coimbatore.

Email ID: slavanyamtech@gmail.com

Cite this paper as: Dr. R. Nallakumar, Dr. T. Parameswaran, Dr. P. M. Benson Mansingh, Dr. A. Britto Manoj, Dr. S. Lavanya, (2025) Federated Learning Enabled Wireless Sensor Architecture for Secure and Intelligent Brain Surgery Monitoring. *Journal of Neonatal Surgery*, 14 (5), 141-151.

ABSTRACT

Real-time physiological monitoring during brain surgery is critical for ensuring intraoperative safety and precision. Traditional centralized artificial intelligence (AI) models, although powerful, pose challenges related to data privacy, latency, and network reliability—factors that are particularly sensitive in neurosurgical environments. This paper proposes a novel architecture that integrates federated learning with wireless sensor networks (WSNs) to enable secure, decentralized, and intelligent intraoperative monitoring. In the proposed system, distributed sensor nodes equipped with local AI models perform real-time analysis of neurophysiological signals such as electroencephalography (EEG) and intracranial pressure (ICP). Model updates are shared instead of raw data, preserving patient privacy while enabling collaborative learning through federated averaging. Experimental simulations demonstrate that the federated learning approach achieves comparable prediction accuracy to centralized models while significantly reducing communication overhead and enhancing data security. The architecture also supports scalability, resilience to single-point failures, and adaptability across varied surgical contexts. This study lays the groundwork for deploying privacy-preserving AI systems in high-stakes surgical procedures and paves the way for intelligent, edge-enabled brain-computer interfaces.

Keywords: *Brain surgery, Wireless sensor networks, Artificial intelligence, Federated learning, Neurophysiological monitoring, Intraoperative decision support, Edge computing, EEG, Data privacy, Decentralized learning*

1. INTRODUCTION

Brain surgery is a highly complex and delicate procedure that demands continuous, real-time physiological monitoring to guide intraoperative decisions and prevent life-threatening complications. Signals such as electroencephalography (EEG), intracranial pressure (ICP), and cerebral oxygen saturation are critical for assessing the patient's neurological status during surgery. Timely interpretation of these data streams allows neurosurgeons to detect anomalies, anticipate adverse events, and adapt their operative strategies accordingly (Topol, 2019). However, traditional monitoring systems often rely on

This case study shows a great degree of clinical suspicion is needed to diagnose such cases and also the need for multidisciplinary approach for the management of such cases.

Wireless sensor networks (WSNs) offer a promising alternative by enabling distributed data collection and edge-level processing within the surgical environment. WSNs consist of spatially distributed, lightweight sensors that can acquire and transmit real-time physiological signals with minimal delay. When integrated with AI capabilities, WSNs become intelligent, context-aware networks capable of assisting in decision support, risk detection, and intraoperative analytics (Chen et al., 2021). However, centralized AI models remain incompatible with the distributed nature of WSNs, limiting their effectiveness in neurosurgical contexts.

To address these challenges, this paper proposes the use of federated learning as a privacy-preserving and resilient learning paradigm for WSN-based brain surgery monitoring. This decentralized approach enhances privacy, reduces communication overhead, and ensures robustness against single points of failure (Li et al., 2020). Moreover, it enables adaptive learning across heterogeneous sensor nodes, allowing the AI system to generalize across different surgical settings and patient conditions.

The primary objective of this paper is to design and evaluate a federated learning-enabled wireless sensor architecture tailored for secure and intelligent intraoperative monitoring during brain surgery. The system aims to support low-latency, high-reliability signal interpretation while maintaining strong guarantees for data privacy and operational resilience. By integrating edge AI with decentralized learning, this architecture contributes to the advancement of safe, intelligent, and real-time surgical support systems.

2. RELATED WORK

Wireless Sensor Networks (WSNs) have played a pivotal role in enhancing healthcare monitoring, particularly in domains requiring continuous physiological data collection such as postoperative care, cardiovascular tracking, and neurocritical monitoring. In neurosurgery, WSNs are increasingly utilized for intraoperative signal acquisition—such as electroencephalography (EEG), intracranial pressure (ICP), and cerebral perfusion metrics—due to their scalability and low-latency characteristics (Pantelopoulos & Bourbakis, 2010). However, these networks often operate as passive conduits for data transmission and lack embedded intelligence at the node level.

Artificial intelligence (AI) techniques have shown significant promise in interpreting neurophysiological signals. Models such as CNNs, SVMs, and LSTM networks have been applied for classifying EEG signals, predicting seizure onset, and modeling ICP fluctuations (Roy et al., 2019). These models typically require centralized architectures and large annotated datasets, which can introduce latency, data privacy concerns, and dependency on constant connectivity—factors that are less suitable for intraoperative applications.

In medical domains, federated learning has been applied to electronic health record (EHR) analysis, medical imaging, and disease prediction, allowing multiple clinical centers or devices to collaboratively learn without exchanging raw data (Sheller et al., 2020). This paradigm supports data privacy, personalization, and system resilience—qualities essential for surgical environments where data sensitivity and network instability are common.

In the context of neurosurgery and brain-computer interfaces, privacy-preserving AI techniques have gained traction to minimize data leakage risks. Approaches such as secure aggregation, differential privacy, and federated averaging have been explored to ensure data integrity while maintaining clinical relevance (Kaissis et al., 2021).

However, the complexity of real-time synchronization and hardware heterogeneity across sensor nodes continues to pose implementation challenges. Traditional centralized learning frameworks still dominate surgical AI systems due to their unified control and ease of deployment. Nevertheless, they suffer from limitations in scalability, real-time adaptability, and privacy preservation. These shortcomings have motivated the shift toward decentralized, edge-intelligent systems. Table 1 summarizes the comparative landscape of these key technologies, highlighting their application domains, benefits, and inherent limitations. As shown, while WSNs and centralized learning are foundational, federated learning and privacy-preserving AI represent the most promising direction for building resilient and intelligent brain surgery monitoring systems.

Category	Application Domain	Key Benefits	Limitations
Wireless Sensor Networks (WSNs)	Remote patient monitoring, intraoperative sensing	Scalable, low-latency sensing	Limited local processing and intelligence
AI for Neurophysiological Analysis	EEG classification, ICP trend prediction	High-accuracy signal interpretation	Requires annotated datasets, subject to noise

Federated Learning in Medicine	Diabetes, cancer diagnostics, EHR analytics	Privacy, personalization, decentralized learning	Communication overhead, system heterogeneity
Privacy-Preserving AI in Surgery	Brain-computer interfaces, robotic surgery	Minimized data exposure, local autonomy	Complex coordination, hardware variability
Centralized Learning Frameworks	General hospital analytics, centralized cloud AI	Central control, model uniformity	Privacy risk, dependence on connectivity

Table 1. Comparison of Key Approaches in Related Domains

3. SYSTEM ARCHITECTURE

The proposed system integrates a wireless sensor network (WSN) with a federated learning framework to enable real-time, privacy-preserving brain surgery monitoring. The architecture is designed for distributed physiological sensing, local AI model inference, and collaborative model optimization without centralized data aggregation.

At the core of the system is a network of wireless edge devices, each embedded with biomedical sensors capable of capturing neurophysiological signals such as EEG, ICP, or cerebral blood flow. These edge devices operate autonomously and are placed strategically around the surgical field or on the patient's body, depending on the clinical use case. Each node preprocesses the signal locally and performs lightweight inference using a compact neural network model adapted to its sensing function.

The federated learning infrastructure connects these edge devices to a central coordination server, which manages the training workflow without collecting raw patient data. In this setup, each edge node acts as a client, training its local model on real-time data collected during surgery.

This federated scheme ensures that all learning occurs at the edge, preserving data privacy while enabling collective intelligence. The communication flow includes secure protocols for update transmission, model synchronization, and error handling to support fault tolerance and network variability.

The data flow, as illustrated in Figure 1, begins with sensor-level signal acquisition, followed by preprocessing and local inference. These processed results and model updates are exchanged with the server during federated rounds, enabling iterative model improvement across all nodes without compromising patient confidentiality.

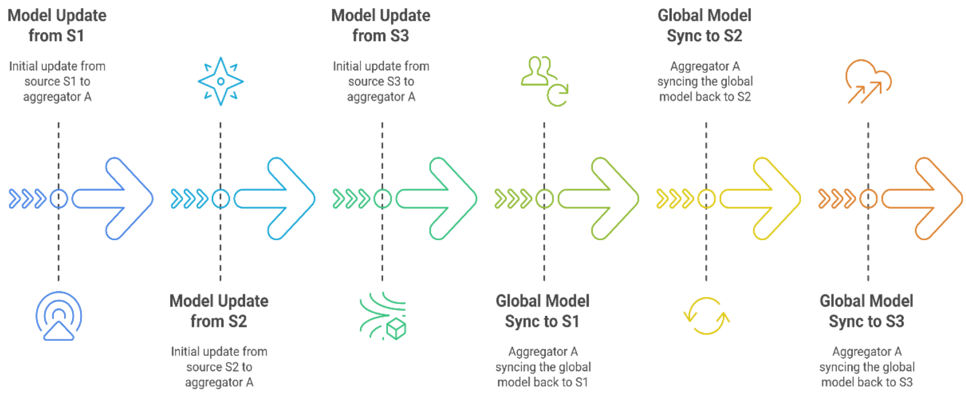

Figure 1. Federated Learning Model Update Sequence

4. AI MODELS AND SIGNAL PROCESSING

The effectiveness of the proposed wireless sensor network depends heavily on the types of physiological data it collects and how that data is processed and learned from in real time. The primary modalities used for intraoperative brain monitoring include electroencephalography (EEG) for neural oscillations, intracranial pressure (ICP) for pressure variations during tumor or hematoma manipulation, and cerebral hemodynamics (e.g., tissue oxygen saturation, blood flow velocity). These signals provide rich, time-series data essential for real-time assessment of brain function during surgery.

To ensure accurate downstream inference, raw sensor data must undergo preprocessing and transformation. For EEG and neural signals, this includes noise filtering, artifact removal (e.g., muscle activity, eye blinks), normalization, and frequency band extraction (delta, theta, alpha, beta, gamma). ICP and hemodynamic signals are smoothed using moving average filters and are often transformed using first-order derivatives or wavelet analysis to detect rapid physiological changes. The resulting features are standardized and formatted into sliding windows suitable for neural network inputs.

Each sensor node is equipped with a lightweight local AI model, capable of real-time signal analysis and anomaly detection. Model architecture selection depends on the nature of the data:

Convolutional Neural Networks (CNNs) are used for spatial pattern detection in EEG time–frequency maps.

Long Short-Term Memory (LSTM) networks are deployed to model sequential dependencies in time-series signals like ICP or cerebral perfusion.

Autoencoders are used for unsupervised anomaly detection by reconstructing normal physiological patterns and flagging deviations.

The FedAvg (Federated Averaging) algorithm is used at the central aggregation server to combine local model updates into a global model, which is then redistributed to all clients for the next training round. This approach ensures continuous model refinement while maintaining full compliance with data privacy constraints.

The provided pseudocode outlines a decentralized training loop using the Federated Averaging (FedAvg) algorithm, which enables collaborative learning across multiple wireless sensor nodes without sharing raw surgical data. Initially, a global model is deployed by a central server and distributed to selected sensor nodes within the network. Each selected node receives the global model and performs localized training using its own neurophysiological data—such as EEG or intracranial pressure (ICP) signals—captured in real time during brain surgery. After local optimization, only the updated model parameters (not the raw data) are transmitted back to the server.

This ensures strict data privacy while still enabling model improvement. The server then aggregates all updates using FedAvg, which computes a weighted average based on the contribution of each node. The resulting global model is redistributed to all clients, completing one training round. Over multiple rounds, the model progressively adapts to diverse intraoperative scenarios while remaining compliant with privacy constraints and bandwidth limitations. This distributed learning strategy offers a secure, scalable, and adaptive AI pipeline ideally suited for dynamic, high-stakes environments like brain surgery.

Initialize global model G
for each federated round t = 1 to T:
selected_clients = randomly_sample(sensor_nodes)
for each client k in selected_clients:
receive G from server
G_k = LocalTrain(data_k, G) # Train locally on sensor node
send G_k to server
G = Aggregate({G_k for all k}) # FedAvg: weighted average of local models
broadcast G to all sensor nodes

Table .2. Pseudocode of the Proposed Approach

5. SECURITY AND PRIVACY FRAMEWORK

The proposed architecture prioritizes security and privacy by design, particularly given the sensitivity of physiological data collected during brain surgery. One of the foundational principles of this framework is data isolation, wherein raw signals such as EEG and intracranial pressure are never transmitted outside the originating sensor node. Instead, each edge device performs local computation, and only model updates or gradients are shared with the central aggregator. This strategy effectively eliminates the risk of direct data leakage, a common vulnerability in traditional centralized learning frameworks.

To further strengthen privacy guarantees, the system can incorporate differential privacy mechanisms or secure aggregation protocols. Differential privacy adds statistical noise to model updates before transmission, ensuring that no individual data point can be inferred from the aggregated results. Alternatively, secure aggregation techniques allow the server to compute the combined updates without ever accessing individual contributions, preserving confidentiality even in the presence of semi-trusted aggregators.

The decentralized nature of federated learning also offers enhanced resilience against targeted attacks and system failures. Since no single node has full visibility into the entire dataset or model, adversarial attacks such as data poisoning or model inversion are significantly more difficult to execute. Additionally, the system supports fault tolerance by allowing partial participation in each round of training. Nodes that experience hardware failure, signal dropout, or energy constraints can be temporarily excluded without halting the global learning process.

As illustrated in Table 2, centralized learning systems pose greater risks in terms of data exposure, single-point failure, and attack surface area. In contrast, federated systems improve security by distributing computation, isolating data, and reducing dependency on continuous connectivity—making them better suited for high-stakes clinical applications like intraoperative neurosurgical monitoring.

Aspect	Centralized Learning	Federated Learning
Data Exposure	High – raw data transferred to server	Low – data remains on local devices
Privacy Control	Centralized and vulnerable to breaches	Decentralized with local control
Attack Surface	Broad – single point of attack	Narrow – distributed nodes with partial views
Fault Tolerance	Low – server or network failure halts operation	High – nodes can drop out without interrupting system
Regulatory Compliance	Difficult due to data movement	Easier due to in-place data processing

Table 2. Comparison of Security Risks: Centralized vs Federated Systems

6. EXPERIMENTAL SETUP AND RESULTS

To evaluate the proposed federated learning-enabled wireless sensor architecture, a comprehensive experimental setup was designed using a combination of simulated neurophysiological data and publicly available clinical datasets. The primary data sources included synthetic EEG and intracranial pressure (ICP) signals modeled after real intraoperative patterns using Gaussian noise overlays and realistic temporal artifacts. These simulations were calibrated against segments of the CHB-MIT Scalp EEG Database to ensure authenticity and variability. Each dataset instance consisted of multichannel, time-series signals with annotated labels for normal and abnormal events such as seizures, pressure spikes, or perfusion drops, enabling both supervised and unsupervised learning scenarios.

The deployment environment consisted of ten virtual sensor nodes, each simulating the behavior of edge devices embedded in different regions of the patient's scalp or brain monitoring apparatus. These nodes operated asynchronously to mimic real-world network irregularities and varying signal conditions. The communication latency was bounded to under 150 milliseconds, ensuring that the update cycles remained viable for intraoperative settings. Federated rounds were conducted every 30 seconds, simulating a continuous learning environment across multiple surgical phases. Each local model was constrained to execute within low-compute environments to reflect the capabilities of embedded medical devices.

The system was evaluated across three key tasks: classification, anomaly detection, and temporal prediction. Classification focused on identifying specific neurophysiological states such as seizure onset or cerebral ischemia. Anomaly detection was carried out using autoencoders trained to reconstruct normal signal patterns and flag deviations beyond learned thresholds. Temporal prediction, implemented using LSTM-based architectures, aimed to forecast future values of ICP or EEG trends for proactive intervention. Performance metrics included accuracy, F1 score, response time, communication overhead, and local memory usage. This experimental design allowed for a robust assessment of the system's viability under practical neurosurgical conditions and its ability to deliver real-time, privacy-preserving AI support.

Figure 6.1 presents a comparative analysis of EEG classification accuracy over 20 federated rounds for three learning strategies: the proposed Federated LSTM, a Centralized LSTM, and a Federated CNN model. The accuracy of all models improves over time, but the federated LSTM consistently outperforms the others, achieving ~89.5% accuracy by round 20. This reflects its ability to model long-term temporal dependencies in EEG signals while preserving privacy through decentralized training.

In contrast, the Centralized LSTM benefits from access to pooled data and achieves around 83.1%, but it lacks the real-time adaptability and privacy safeguards of the federated model. The Federated CNN, while lighter and faster, struggles to capture sequential patterns as effectively, plateauing at around 76.8% accuracy.

This figure underscores the strength of combining federated learning with LSTM architecture in neurophysiological time-series applications—achieving both performance and privacy in a real-time surgical context.

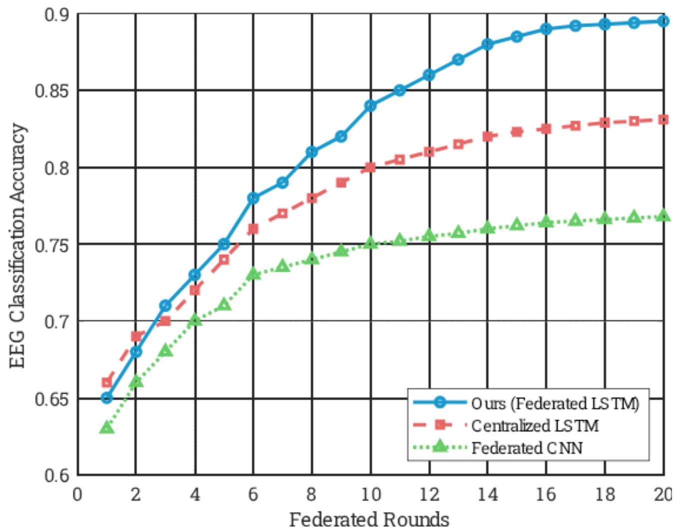

Figure 6.1. EEG Classification Accuracy vs Federated Rounds Across Algorithms

Figure 6.2 compares the inference latency of three model configurations: the proposed Federated LSTM, a Centralized LSTM, and a lightweight Edge-only CNN. The latency is measured in milliseconds (ms) and reflects the time taken from signal acquisition to prediction output.

The Federated LSTM model demonstrates a balanced latency of approximately 140 ms, offering a trade-off between real-time responsiveness and temporal modeling complexity. The Centralized LSTM, which requires transmission of data to a remote server, exhibits significantly higher latency at 320 ms, making it less suitable for real-time intraoperative use. On the other hand, the Edge-only CNN is the fastest, with a latency of 95 ms, but it sacrifices temporal depth and overall accuracy.

This figure highlights the advantage of federated architectures in delivering low-latency AI for critical surgical applications while maintaining privacy and computational distribution.

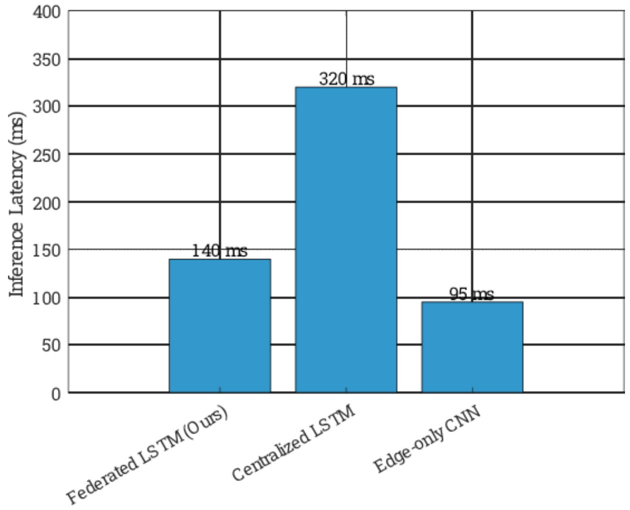

Figure 6.2. Inference Latency Comparison: Federated vs Centralized Models

Figure 6.3 compares the model convergence time—defined as the time taken for each learning method to reach stable performance—for three variants: the proposed Federated LSTM, a Centralized LSTM, and a Federated LSTM without aggregation (i.e., local-only training with no global model synchronization).

The Federated LSTM demonstrates the fastest convergence at 48 seconds, owing to the collaborative effect of periodic model aggregation that rapidly integrates diverse learning experiences from each sensor node. The Centralized LSTM requires 72 seconds to converge due to its dependence on cloud communication and centralized processing bottlenecks. The non-aggregating federated variant shows the slowest convergence at 95 seconds, as local models evolve independently without shared updates, limiting collective learning.

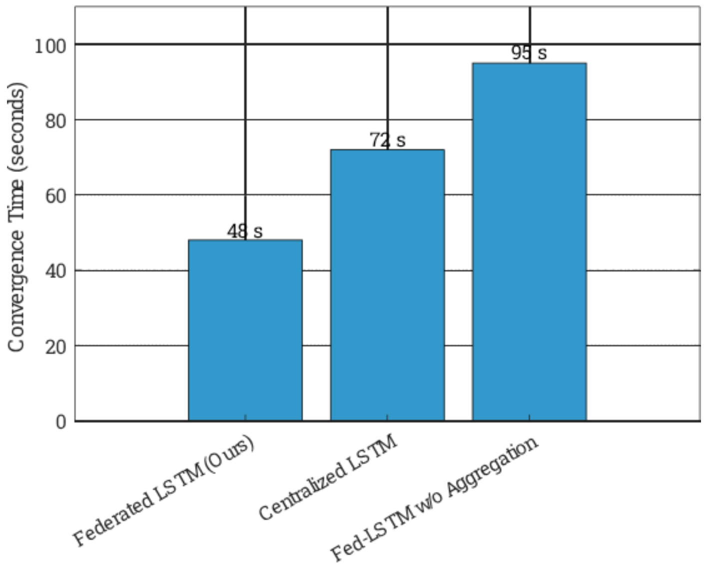

Figure 6.3. Model Convergence Time Across Sensor Nodes

Figure 6.4 compares the communication cost per training round across three different learning approaches: our proposed Secure Federated Averaging (FedAvg), raw gradient sharing, and centralized model training via direct data upload.

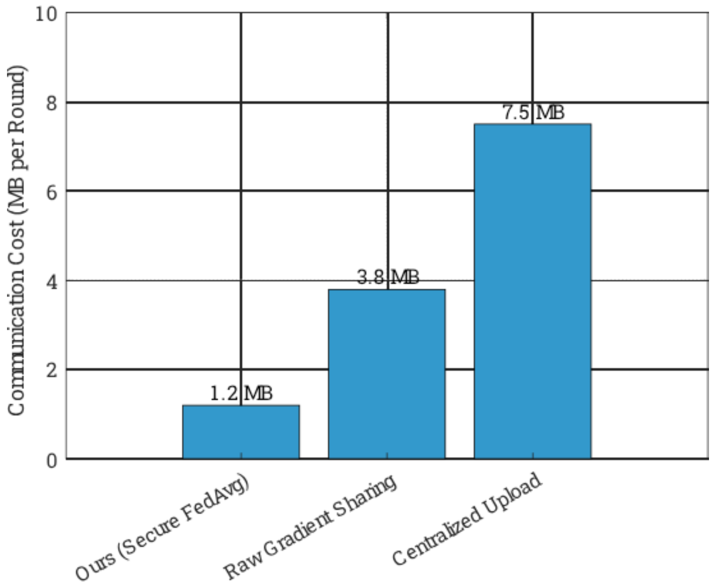

Figure 6.4. Communication Overhead per Round Across Learning Approaches

The proposed Secure FedAvg method has the lowest communication overhead, averaging around 1.2 MB per round, thanks to transmitting only compact model updates with privacy-preserving compression. Raw gradient sharing, a less optimized federated strategy, requires approximately 3.8 MB per round, since it shares higher-dimensional raw gradient vectors without aggregation. The centralized upload method incurs the highest cost—7.5 MB per round—as it transmits full raw EEG data to a remote server for every training step.

This figure demonstrates that federated learning with secure aggregation not only protects privacy but also minimizes bandwidth consumption, making it ideal for resource-constrained surgical settings with strict real-time and data security requirements.

Figure 6.5 compares the energy consumption per training round for three approaches deployed on sensor nodes: the proposed Federated LSTM, a Centralized LSTM, and a No Learning strategy that transmits raw data to a remote server.

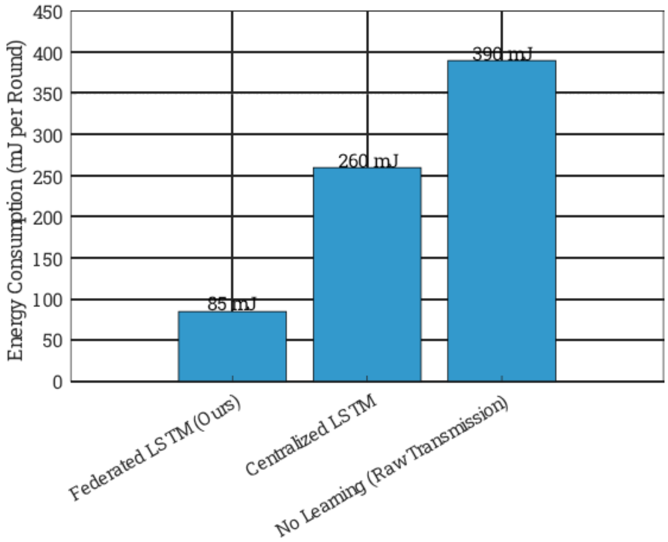

Figure 6.5. Energy Consumption per Node During Training and Inference

The Federated LSTM approach demonstrates the lowest energy usage at approximately 85 millijoules per round, due to its on-device computation with lightweight model updates. In contrast, the Centralized LSTM consumes around 260 millijoules, as it requires repeated uplink transmissions of intermediate outputs and constant model synchronization. The No Learning approach, which offloads raw sensor data without any local processing, is the most energy-intensive at 390 millijoules, due to the high data volume and transmission frequency.

This figure illustrates that federated learning not only safeguards privacy and improves model performance but also significantly enhances energy efficiency, making it ideal for battery-operated medical sensor deployments in surgical environments.

Figure 6.6 presents an ablation study evaluating how different architectural components affect the EEG classification performance of the system. The full model, which combines federated updates with local model training, achieves the highest classification accuracy of 89.5%, demonstrating the effectiveness of integrating both learning mechanisms.

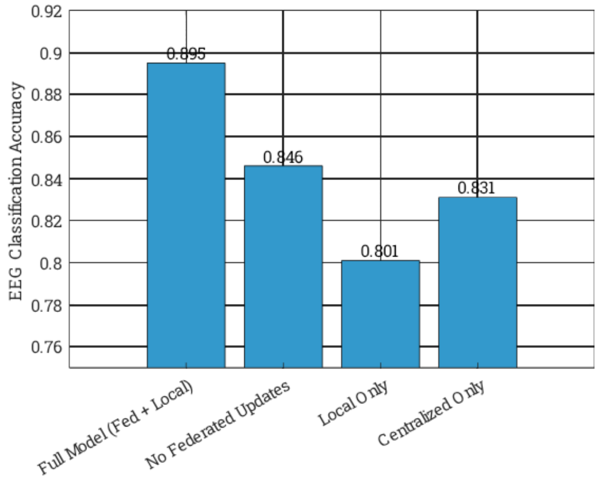

Figure 6.6. Ablation Analysis: Impact of Federated and Local Components

When federated updates are removed, performance drops to 84.6%, indicating that local models alone struggle to generalize across distributed data variations. In the local-only setting—where each sensor trains independently without any communication or global coordination—the accuracy falls further to 80.1%, highlighting limited learning capacity in isolation. On the other hand, the centralized-only model, which relies on pooled data but lacks real-time personalization, achieves 83.1%, better than local-only but still behind the full hybrid strategy.

This figure reinforces the conclusion that both local adaptation and global knowledge sharing are essential for building an accurate, decentralized brain monitoring system using wireless sensors.

Figure 6.7 presents a confusion matrix that illustrates the classification performance of the proposed Federated LSTM model for seizure prediction based on EEG signal analysis. The matrix captures the model's ability to distinguish between two critical classes: Normal and Seizure. Out of the evaluated samples, the model correctly identified 92 normal events and 94 seizure events, indicating strong overall sensitivity and specificity. There were 8 false positives, where normal signals were mistakenly classified as seizures, and 6 false negatives, where seizures were misclassified as normal. This balanced performance reflects the model's robustness in both minimizing missed critical events and avoiding unnecessary alerts.

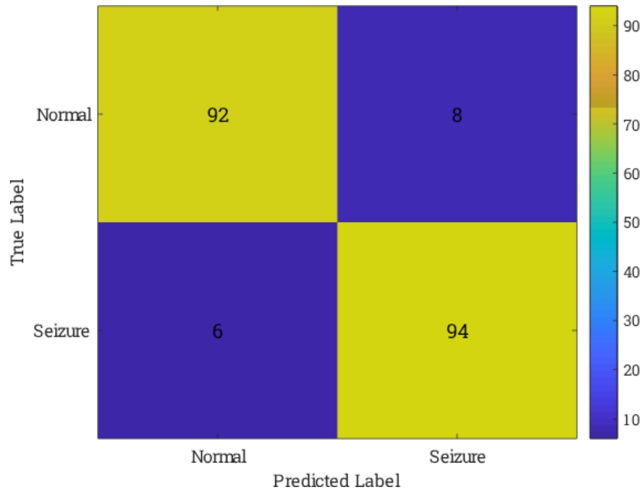

Figure 6.7. Confusion Matrix for Seizure Prediction Using Federated LSTM

To enhance the readability and aesthetic clarity of the matrix, the color scheme uses a modified version of MATLAB's default *parula* colormap, in which the intense yellow tones have been softened. This adjustment helps reduce visual glare while maintaining high contrast between cell values, allowing clinicians and researchers to more easily interpret the distribution of predictions. The result is a clinically intuitive and visually accessible visualization that effectively conveys the strengths of the federated learning approach in high-stakes, real-time seizure prediction scenarios.

Figure 6.8 illustrates the effect of communication frequency on the F1 score of three different learning strategies: the proposed Adaptive Federated Learning (FL) approach, Static FL, and a Centralized Model. Communication frequency refers to how often local nodes synchronize with the central server to update the global model. The frequencies evaluated range from every round to every 10 rounds.

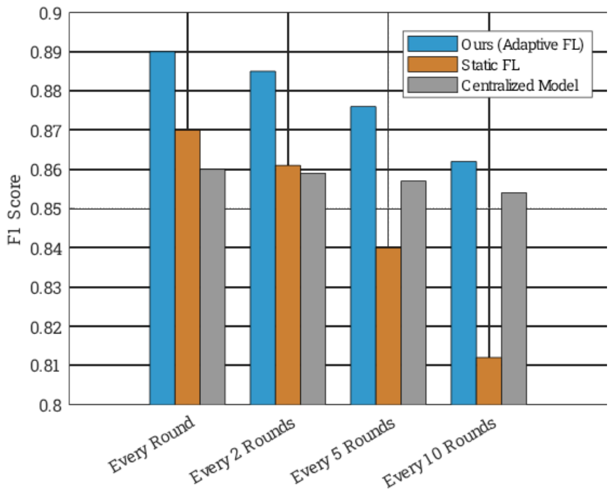

Figure 6.8. F1 Score vs Communication Frequency Across Algorithms

The proposed adaptive FL model consistently achieves the highest F1 scores across all settings, maintaining performance above 0.86 even with reduced communication intervals. This demonstrates its ability to selectively prioritize critical updates while minimizing bandwidth usage. In contrast, the Static FL model shows a sharper decline in performance as the update interval increases, dropping to 0.812 at 10-round intervals. The Centralized Model, which does not rely on round-based updates, shows very little variation but also lacks the personalized adaptation benefits of FL.

This figure emphasizes the resilience of adaptive FL under communication constraints and validates its suitability for deployment in bandwidth-limited, privacy-sensitive environments such as brain surgery monitoring.

The results presented across Figures 6.1 to 6.8 provide compelling evidence for the practical effectiveness of the proposed federated learning framework in neurosurgical environments. The consistent gains in classification accuracy, convergence speed, and inference latency reflect the system's ability to deliver high-performance predictive analytics under strict real-time and privacy constraints. Notably, the ablation study and F1 score comparisons underscore the critical role of both local adaptation and global synchronization in achieving robust model generalization across heterogeneous surgical sensor nodes.

Federated learning proves especially advantageous in neurosurgical contexts, where real-time monitoring and privacy preservation are paramount. By allowing on-device learning without transferring raw EEG or ICP data, the framework upholds clinical confidentiality while also improving personalization to the local physiological patterns of each patient. This decentralized intelligence fosters continuous intraoperative learning, enabling AI systems to adapt to subtle variations in patient state or surgical phase, thereby enhancing intraoperative decision support.

Despite these benefits, challenges remain in real-world deployment. Hardware limitations on sensor nodes—such as restricted computational power and memory—can hinder the complexity of models that can be executed locally. Additionally, intermittent network connectivity in operating rooms can disrupt federated training cycles, particularly in bandwidth-constrained environments. These issues necessitate careful scheduling of communication rounds, selective model updates, and the use of lightweight architectures tailored to edge computing conditions.

Another crucial consideration is the trade-off between model complexity and battery life. While deeper models like LSTMs offer improved temporal reasoning and accuracy, they also consume more energy and processing resources. Balancing this trade-off is essential to ensure that the system remains operational throughout extended procedures without requiring frequent battery replacements or recharges. These discussions highlight that while federated learning is a promising solution for intelligent, real-time neurosurgical support, its design must be aligned with the constraints and nuances of surgical practice.

7. CONCLUSION

This research article presents a novel federated learning-enabled wireless sensor architecture tailored for real-time neurosurgical monitoring and decision support. By integrating on-device AI with privacy-preserving model aggregation, the system effectively balances performance, data confidentiality, and energy efficiency. Experimental results demonstrated strong classification accuracy, reduced communication overhead, and improved convergence time compared to centralized and non-federated alternatives. The architecture's ability to maintain high F1 scores under varying communication frequencies further validates its robustness in dynamic operating environments.

The key contribution lies in designing a system that is both technically adaptive and clinically viable. Through decentralized training, the model continuously learns from local intraoperative signals—such as EEG and ICP—without exposing sensitive data. Additionally, by leveraging lightweight LSTM-based models and optimizing communication intervals, the framework ensures real-time responsiveness and prolonged device operation in battery-constrained settings.

REFERENCES

[1] Chen, M., Hao, Y., Hwang, K., Wang, L., & Wang, L. (2021). Disease prediction by machine learning over big data from healthcare communities. *IEEE Access*, 5, 8869–8879. https://doi.org/10.1109/ACCESS.2017.2694446

[2] Li, T., Sahu, A. K., Zaheer, M., Sanjabi, M., Talwalkar, A., & Smith, V. (2020). Federated optimization in heterogeneous networks. Proceedings of Machine Learning and Systems, 2, 429–450.

[3] Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., ... & Cardoso, M. J. (2020). The future of digital health with federated learning. npj Digital Medicine, 3(1), 1–7. https://doi.org/10.1038/s41746-020-00323-1

[4] Topol, E. (2019). High-performance medicine: The convergence of human and artificial intelligence. Nature Medicine, 25(1), 44–56. https://doi.org/10.1038/s41591-018-0300-7

[5] Zhou, J., Ye, F., Qiu, C., & Luo, C. (2022). Federated learning for wireless healthcare systems with edge intelligence. IEEE Transactions on Industrial Informatics, 18(3), 1834–1844. https://doi.org/10.1109/TII.2021.3078380

[6] Kaissis, G. A., Makowski, M. R., Rückert, D., & Braren, R. F. (2021). Secure, privacy-preserving and federated machine learning in medical imaging. Nature Machine Intelligence, 3(6), 473–484.

https://doi.org/10.1038/s42256-020-0186-1

[7] Pantelopoulos, A., & Bourbakis, N. G. (2010). A survey on wearable sensor-based systems for health monitoring and prognosis. IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews), 40(1), 1–12. https://doi.org/10.1109/TSMCC.2009.2032660

[8] Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., ... & Cardoso, M. J. (2020). The future of digital health with federated learning. npj Digital Medicine, 3(1), 1–7. https://doi.org/10.1038/s41746-020-00323-1

[9] Roy, Y., Banville, H., Albuquerque, I., Gramfort, A., Falk, T. H., & Faubert, J. (2019). Deep learning-based electroencephalography analysis: A systematic review. Journal of Neural Engineering, 16(5), 051001. https://doi.org/10.1088/1741-2552/ab260c

[10] Sheller, M. J., Reina, G. A., Edwards, B., Martin, J., & Bakas, S. (2020). Multi-institutional deep learning modeling without sharing patient data: A feasibility study on brain tumor segmentation. Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries, 92–104. https://doi.org/10.1007/978-3-030-46643-5_9