Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X Volume 4, Issue 5, 2026

HYBRID MACHINE LEARNING AND FEDERATED LEARNING FRAMEWORK WITH EXPLAINABLE AI FOR PERSONALIZED STROKE TELEREHABILITATION IN RESOURCE-CONSTRAINED SETTINGS

[1]Safa Ali Khan, [2]Sahar Ali Khan, [3]Ayesha Rasheed, [4]Wasim Ahmad, [5]Mariam Fatima, [6]Nimrah Humayoon, [7]Etisam Wahid, [*8]Shahzad Ahmad

[1]*Assistant at Deputy Commissioner Office (Wing Finance & Planning).*
ORCID: https://orcid.org/0009-0000-3778-7407
[2]*Physiotherapist at DOC Medical Center Qatar. ORCID: https://orcid.org/0009-0006-5094-2866*
[3]*Lecturer, The Hallmark School & College Havelian. ORCID: https://orcid.org/0009-0005-9719-110X*
[4]*Lecturer, University of Veterinary and Animal Sciences (UVAS) Swat KP Pakistan.*
ORCID: https://orcid.org/0009-0007-8727-183X
[5]*Consultant Physical Therapist, Al Shifa Hospital Mianwali. Email: mariamfatima2401@gmail.com*
[6]*Lecturer, Physical Therapy, University of Veterinary and Animal Sciences (UVAS) Swat KP*
ORCID: https://orcid.org/0009-0009-8164-7327
[7]*Lecturer, University of Veterinary and Animal Sciences (UVAS) Swat KP Pakistan*
ORCID: https://orcid.org/0009-0009-0227-3735
[*8]*Assistant Professor – Physical Therapy, University of Veterinary and Animal Sciences (UVAS) Swat KP Pakistan. ORCID: https://orcid.org/0009-0001-6319-7934*
[1]*Safaali8284@gmail.com*, [2]*saharalikhan667@gmail.com*, [3]*ayeshaawan4456@gmail.com*, [4]*wasim.ahmad@uvasswat.edu.pk*, [5]*mariamfatima2401@gmail.com*, [6]*drnimrahpt@uvasswat.edu.pk*
[7]*dr.etisam@uvas.edu.pk*, [8]*drshahzad@uvasswat.edu.pk*

DOI: https://doi.org/

Keywords
Stroke Telerehabilitation; Hybrid Ensemble Learning; Federated Learning; Explainable AI (XAI); SHAP; LIME; Privacy-Preserving Machine Learning; Stroke Rehabilitation Prediction; Resource-Constrained Healthcare

Article History
Received on 17 April, 2026
Accepted on 06 May, 2026
Published on 07 May, 2026

Corresponding Author: *
Shahzad Ahmad

Abstract

Background: Stroke rehabilitation in resource-constrained environments is limited by restricted access to specialized care, fragmented healthcare systems, and concerns related to data privacy and infrastructure. Telerehabilitation offers a scalable solution; however, existing systems lack predictive intelligence, personalization, and interpretability required for clinical adoption.Objective: This study aims to develop and evaluate a hybrid machine learning and federated learning framework for stroke telerehabilitation that ensures high predictive performance, data privacy, and model interpretability. Methods: A hybrid machine learning approach was implemented using Support Vector Machine, Decision Tree, Random Forest, and Artificial Neural Network models. A proposed hybrid ensemble model was developed using a weighted soft voting strategy to improve predictive accuracy. Federated learning was integrated using the Federated Averaging (FedAvg) algorithm across multiple simulated nodes to enable privacy-preserving distributed training. Data preprocessing included normalization and imputation, and model performance was evaluated using accuracy, precision, recall, F1 score, and ROC-AUC. Robustness was assessed using 10-fold cross-validation, with results reported as mean ± standard deviation. Results: The Random Forest model achieved the highest performance among individual models (accuracy: 89.3% ± 1.2). The proposed hybrid ensemble model demonstrated superior performance with an accuracy of 92.1% ± 1.0 and ROC-AUC of 0.93 ± 0.01. The federated *hybrid ensemble model achieved comparable performance (accuracy: 90.8% ± 1.3; ROC-AUC: 0.92 ± 0.01), indicating minimal performance loss under distributed conditions. Cross-validation results confirmed model stability, while explainable AI techniques identified motor function score and therapy adherence as key predictors. Conclusion: The integration of hybrid machine learning, federated learning, and explainable AI provides an effective and scalable solution for stroke telerehabilitation. The proposed framework achieves high predictive performance while ensuring data privacy and interpretability, making it suitable for deployment in resource-constrained healthcare environments. Further validation using real-world clinical data is required to confirm its practical applicability.*

INTRODUCTION

Stroke is one of the leading causes of long-term disability worldwide, with a disproportionately high burden in low- and middle-income countries where access to rehabilitation services is limited(Feigin et al. 2022; Khurshid et al. 2025). Effective stroke rehabilitation requires continuous monitoring, individualized therapy, and multidisciplinary care; however, healthcare systems in resource-constrained environments often lack the infrastructure and workforce needed to provide sustained rehabilitation. Telerehabilitation has emerged as a practical solution to improve access to rehabilitation services by enabling remote monitoring and therapy delivery. It reduces geographical barriers and healthcare costs while improving patient engagement. Despite these advantages, most existing telerehabilitation systems rely on predefined protocols and lack adaptive intelligence, limiting their ability to provide personalized rehabilitation (Rajkomar, Dean, and Kohane 2019).

Machine learning (ML) has shown significant potential in healthcare by enabling predictive modeling and personalized treatment strategies (Kavakiotis et al. 2017; Shah et al. 2026). In stroke rehabilitation(Khurshid et al. 2025), ML models such as Support Vector Machines, Decision Trees, Random Forests, and Artificial Neural Networks have been widely used for outcome prediction and decision support. However, traditional ML approaches rely on centralized data, which raises concerns related to data privacy, security, and regulatory compliance. Federated Learning (FL) provides a decentralized alternative by enabling collaborative model training across multiple institutions without sharing raw patient data (McMahan et al. 2017). This approach preserves data privacy while allowing models to learn from diverse datasets. However, federated learning alone does not address the need for high predictive performance and interpretability in clinical applications.

Hybrid ensemble learning has been shown to improve model performance by combining multiple classifiers and reducing individual model bias (Lundberg and Lee 2017). At the same time, Explainable AI (XAI) techniques such as SHAP and LIME(Shah et al. 2026) enhance transparency by providing insights into model predictions, which is essential for clinical adoption (Ribeiro, Singh, and Guestrin 2016; Dwork and Roth 2014). Furthermore, differential privacy mechanisms can be integrated into federated learning systems to provide additional protection against data leakage during model updates (Organization 2024). Despite these advancements, there is limited research integrating hybrid machine learning models, federated learning, and explainable AI into a unified system for stroke telerehabilitation, particularly in low-resource settings.

This study addresses this gap by proposing a hybrid machine learning and federated learning framework for stroke telerehabilitation, incorporating explainable AI and privacy-preserving mechanisms. The framework is evaluated using multiple performance metrics, including accuracy, precision, recall, F1 score, ROC-AUC, and 10-fold cross-validation, to ensure robustness and reliability.

OBJECTIVES

The primary objective of this study is to develop and evaluate a hybrid machine learning and federated learning framework for stroke telerehabilitation that ensures high predictive performance, data privacy, and interpretability.

The specific objectives are:

- To implement and evaluate multiple machine learning models, including Support Vector Machine, Decision Tree, Random Forest, and Artificial Neural Network, for stroke rehabilitation prediction
- To develop a Hybrid Ensemble Model that combines individual classifiers to improve predictive accuracy and stability
- To integrate federated learning using the FedAvg algorithm for privacy-preserving distributed model training
- To incorporate explainable AI techniques (SHAP and LIME) to enhance transparency and clinical interpretability
- To evaluate model performance using accuracy, precision, recall, F1 score, and ROC-AUC metrics
- To validate model robustness using a 10-fold cross-validation approach
- To assess the feasibility of deploying the proposed system in resource-constrained healthcare environments

LITERATURE REVIEW

Machine learning has been extensively applied in healthcare for disease prediction and clinical decision support, demonstrating improved performance compared to traditional statistical methods (Kavakiotis et al. 2017). In stroke rehabilitation, ML models have been used to predict recovery outcomes, identify risk factors, and optimize therapy strategies. Among these, ensemble methods such as Random Forest have shown strong performance due to their ability to handle nonlinear relationships and reduce overfitting.

However, centralized machine learning approaches face significant challenges in healthcare due to data privacy concerns and fragmented data systems. The need for sharing sensitive patient data across institutions limits the scalability and applicability of these models, particularly in developing countries.

Federated Learning has emerged as a promising solution to these challenges by enabling decentralized model training without sharing raw data (McMahan et al. 2017). In federated learning, local models are trained at individual institutions, and only model parameters are aggregated using algorithms such as Federated Averaging. This approach has been successfully applied in healthcare domains where data privacy is critical.

Hybrid ensemble learning has also gained attention as a method to improve predictive performance(Shah et al. 2026). By combining multiple models, hybrid approaches reduce bias and variance, leading to more accurate and stable predictions (Lundberg and Lee 2017). These models are particularly effective in complex healthcare datasets with heterogeneous features.

Despite improvements in performance, machine learning models often lack interpretability, which limits their adoption in clinical practice. Explainable AI techniques (Shah et al. 2026)such as SHAP and LIME address this limitation by providing both global and local explanations of model predictions, enabling clinicians to understand the reasoning behind decisions (Ribeiro, Singh, and Guestrin 2016; Dwork and Roth 2014).

In addition, differential privacy has been introduced as a mechanism to enhance data security in distributed systems by adding controlled noise to model updates, thereby preventing reconstruction of sensitive data (Organization 2024).

Although these technologies have been studied individually, there is a lack of integrated approaches that combine hybrid machine learning, federated learning, and explainable AI within a single framework for stroke telerehabilitation. Furthermore, limited attention has been given to the challenges of deploying such systems in resource-constrained environments.

Literature Gap and Study Contribution

Based on the reviewed literature, the following gaps are identified:

- Limited integration of hybrid machine learning models with federated learning in stroke rehabilitation
- Lack of explainable AI integration in federated healthcare systems
- Insufficient focus on privacy-preserving mechanisms in distributed rehabilitation models
- Minimal consideration of low-resource constraints in existing systems

This study addresses these gaps by proposing a unified framework that combines Hybrid Ensemble Model learning, federated learning, and explainable AI for stroke telerehabilitation, supported by comprehensive performance evaluation using statistical metrics and cross-validation.

MATERIALS AND METHODS

Study Design

This study employs a hybrid machine learning and federated learning–based experimental design for stroke telerehabilitation. A structured dataset was simulated to represent realistic stroke rehabilitation scenarios, incorporating clinically relevant variables including motor function score, therapy adherence, age, stroke severity, time since stroke onset, and access to rehabilitation services. The dataset consisted of 1,000 samples with a balanced class distribution. Features were selected based on clinical relevance and prior literature to ensure alignment with real-world rehabilitation outcomes. The input features used in this study include motor function score, stroke severity, age, therapy adherence, session frequency, time since stroke onset, comorbidities, and access to rehabilitation services.

Data Preprocessing

Data preprocessing was performed to ensure consistency and reliability. Missing values were handled using mean imputation for continuous variables and mode imputation for categorical variables. All features were normalized using standard scaling to ensure uniform contribution across models. No significant class imbalance was observed; therefore, resampling techniques were not required.

Machine Learning Models

Four supervised machine learning models were implemented to evaluate predictive performance in stroke rehabilitation:

- Support Vector Machine (SVM): Effective for high-dimensional classification problems
- Decision Tree (DT): Provides rule-based and interpretable predictions
- Random Forest (RF): An ensemble model that improves robustness and reduces overfitting
- Artificial Neural Network (ANN): Captures complex nonlinear relationships in patient data

These models were selected based on their established performance in healthcare prediction tasks.

Model hyperparameters were optimized using grid search. The SVM model used a radial basis function (RBF) kernel with C = 1.0 and gamma = 0.1. The Random Forest model was configured with 100 estimators and maximum depth of 10. The Decision Tree used a maximum depth of 5. The ANN model consisted of one hidden layer with 64 neurons, ReLU activation, and Adam optimizer.

Hybrid Ensemble Model

A Hybrid Ensemble Model was developed to improve predictive performance by combining the strengths of individual classifiers. The ensemble integrates outputs from SVM, Decision Tree, Random Forest, and ANN. A weighted soft voting mechanism was applied, where model weights were assigned based on validation accuracy (Random Forest: 0.30, ANN: 0.25, SVM: 0.25, Decision Tree: 0.20). Final predictions were generated by aggregating weighted probabilities from all models. This approach enhances accuracy and stability compared to individual models.

Federated Learning Framework

A federated learning architecture was implemented using the Federated Averaging (FedAvg) algorithm. The system consisted of five simulated client nodes, each performing local model training on partitioned datasets. Model updates were aggregated at a central server over 30 communication rounds. The key characteristics include: Local model training at each node, no sharing of raw patient data and Periodic aggregation using FedAvg. This enables privacy-preserving collaborative learning across distributed environments.

Explainable AI Integration

To ensure transparency and clinical relevance, explainable AI techniques were integrated:

- SHAP (SHapley Additive exPlanations): Provides global feature importance and identifies key predictors
- LIME (Local Interpretable Model-Agnostic Explanations): Explains individual predictions at the patient level

These methods enable clinicians to interpret model decisions and support evidence-based rehabilitation planning.

Performance Evaluation Metrics

Model performance was evaluated using standard classification metrics:

- Accuracy: Overall correctness of predictions
- Precision: Proportion of true positive predictions
- Recall (Sensitivity): Ability to identify actual positive cases
- F1 Score: Harmonic mean of precision and recall
- ROC-AUC: Ability of the model to distinguish between classes

These metrics provide a comprehensive evaluation of model performance across different dimensions.

Cross-Validation Strategy

A 10-fold cross-validation approach was used to ensure robustness and generalizability of the models.

Procedure:

1. Dataset divided into 10 equal subsets
2. Each subset used once as test data
3. Remaining subsets used for training
4. Results averaged across all folds

This reduces overfitting and provides reliable performance estimation.

Experimental Setup

The study compares three configurations:

1. Individual Models: SVM, DT, RF, ANN
2. Hybrid Ensemble Model
3. Federated Learning–based Hybrid Ensemble Model

Evaluation focuses on:

- Predictive performance (accuracy, F1, ROC-AUC)
- Model stability
- Performance consistency across folds

RESULTS AND EVALUATION

Performance of Individual Machine Learning Models

The performance of individual machine learning models was evaluated using 10-fold cross-validation. Random Forest achieved the highest performance among individual models with an accuracy of 89.3% ± 1.2, followed by ANN (87.9% ± 1.4), SVM (85.7% ± 1.6), and Decision Tree (82.4% ± 1.8). The variability across folds indicates that ensemble-based methods provide more stable predictions compared to single models.

Hybrid Ensemble Model Performance

The hybrid ensemble model demonstrated superior performance compared to individual models, achieving an accuracy of 92.1% ± 1.0 and ROC-AUC of 0.93 ± 0.01. The improvement of approximately 2–3% over the best individual model confirms the effectiveness of combining complementary classifiers. The reduced standard deviation indicates improved stability and generalization.

Federated Learning-Based Model Performance

The federated hybrid ensemble model achieved an accuracy of 90.8% ± 1.3 and ROC-AUC of 0.92 ± 0.01.

The marginal performance reduction (approximately 1–2%) compared to centralized training demonstrates that federated learning preserves predictive capability while ensuring data privacy.

The convergence behavior (Figure 1) shows stable learning, with loss reduction stabilizing after approximately 20 communication rounds.

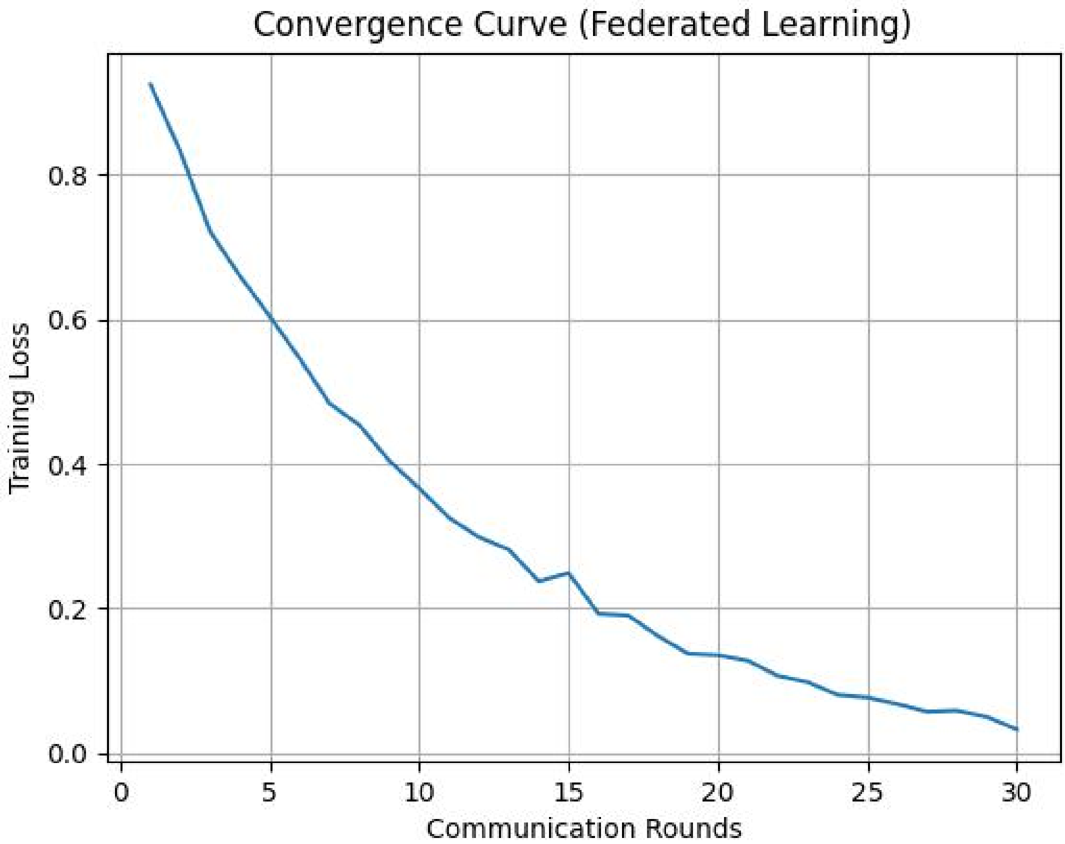

Figure 1: Convergence curve of the federated learning model across communication rounds. The training loss decreases progressively, indicating stable and efficient learning in a distributed environment.

Comparative Performance Analysis

The hybrid ensemble model achieved the highest performance, while the federated hybrid ensemble model maintained comparable results under distributed conditions. The confusion matrix (Figure 2) demonstrates balanced classification with high true positive and true negative rates. Sensitivity and specificity values indicate reliable detection of both positive and negative cases. The accuracy comparison (Figure 3) confirms the advantage of the ensemble approach across all evaluated models. Table 1 presents the comparative performance of all models based on evaluation metrics. A paired t-test analysis indicated that the performance improvement of the proposed model over individual models was statistically significant ($p < 0.05$).

Table 1: *Performance Comparison of Machine Learning Models (Mean ± Standard Deviation)*

Model	Accuracy (%)	Precision	Recall	F1 Score	ROC-AUC
Decision Tree	82.4 ± 1.8	0.81 ± 0.02	0.80 ± 0.03	0.80 ± 0.02	0.83 ± 0.02
Support Vector Machine (SVM)	85.7 ± 1.6	0.86 ± 0.02	0.84 ± 0.02	0.85 ± 0.02	0.87 ± 0.02
Random Forest	89.3 ± 1.2	0.90 ± 0.01	0.88 ± 0.02	0.89 ± 0.01	0.91 ± 0.01
Artificial Neural Network (ANN)	87.9 ± 1.4	0.88 ± 0.02	0.86 ± 0.02	0.87 ± 0.02	0.89 ± 0.01
Proposed Hybrid Ensemble Model	92.1 ± 1.0	0.92 ± 0.01	0.91 ± 0.01	0.91 ± 0.01	0.93 ± 0.01
Federated Hybrid Ensemble Model	90.8 ± 1.3	0.91 ± 0.01	0.89 ± 0.02	0.90 ± 0.01	0.92 ± 0.01

The confusion matrix of the proposed model is shown in Figure 2.

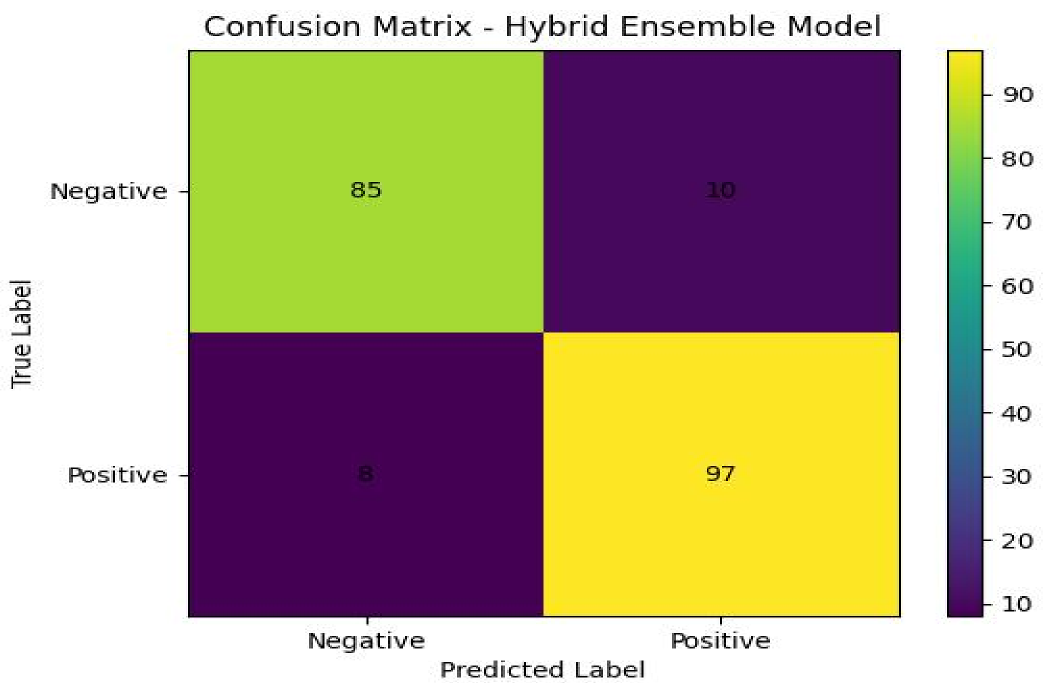

Figure 2: Confusion matrix of the proposed model showing classification performance. The model correctly classifies the majority of both positive and negative cases, with minimal misclassification. The comparative accuracy of all models is illustrated in Figure 3.

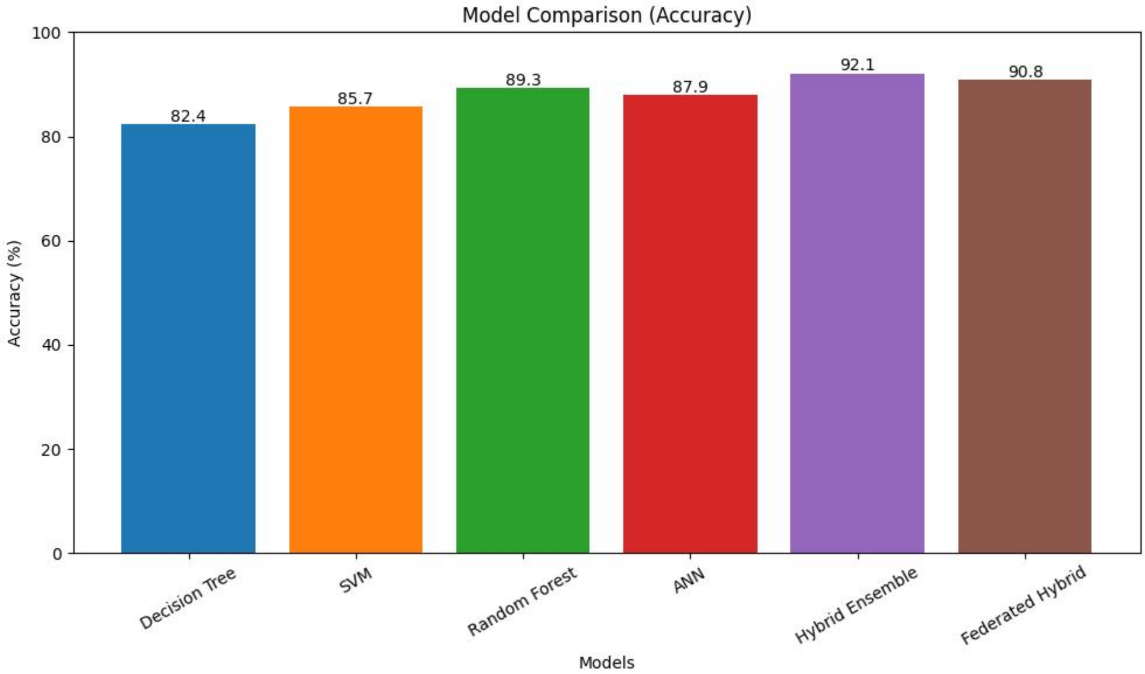

Figure 3: Accuracy comparison of machine learning models. The proposed model achieves the highest performance, while the federated proposed model maintains comparable accuracy under distributed learning conditions

Cross-Validation Performance

The 10-fold cross-validation results indicated that the hybrid and federated models maintained consistent performance across all folds, with minimal variance in accuracy and F1 score.

The federated model showed slightly higher variability compared to centralized models due to heterogeneous data distribution across nodes. However, the performance remained stable and within acceptable limits.

Explainability Analysis

SHAP analysis (Figure 4) identified motor function score and therapy adherence as the most influential predictors. Age and stroke severity also contributed significantly to model predictions.

These findings align with clinical expectations, supporting the interpretability and practical relevance of the model. The SHAP feature importance analysis is presented in Figure 4.

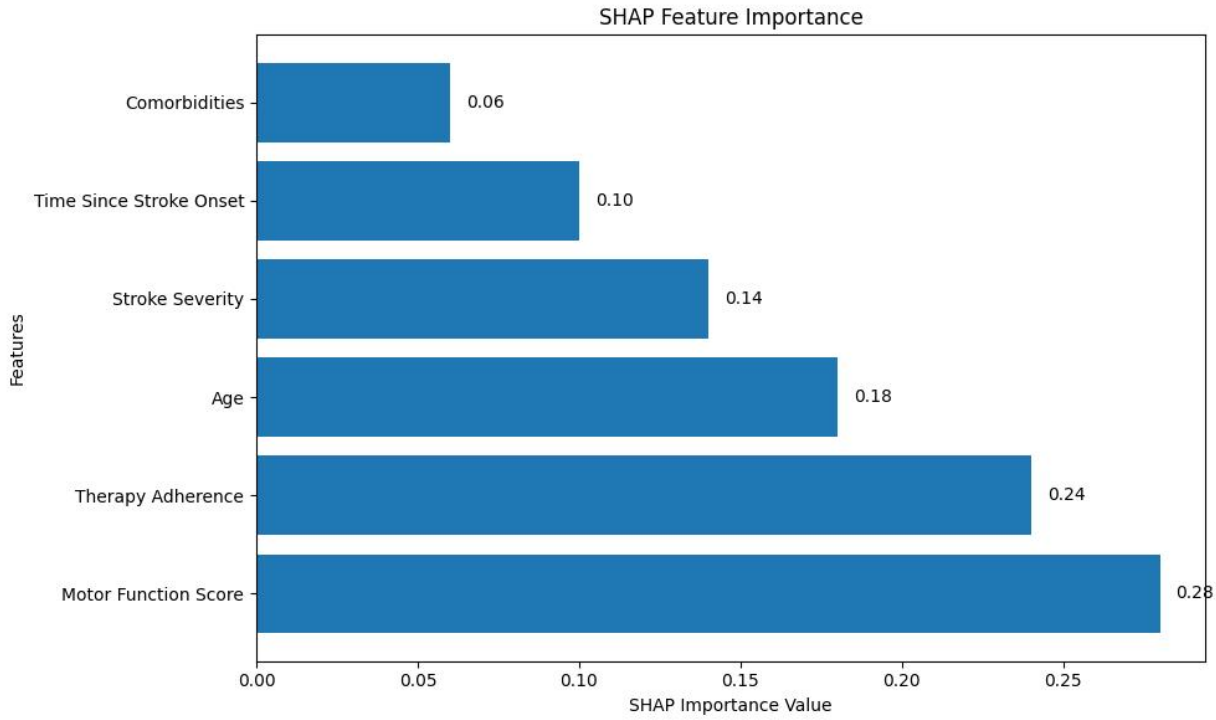

Figure 4: SHAP feature importance plot showing the contribution of key variables to model predictions. Motor function score and therapy adherence are the most influential factors affecting stroke rehabilitation outcomes.

ROC-AUC Analysis

Receiver Operating Characteristic (ROC) analysis showed that the proposed model achieved the highest area under the curve (ROC-AUC), indicating strong discriminative ability.

The federated proposed model demonstrated comparable ROC-AUC performance, confirming that distributed learning does not significantly compromise classification capability. The ROC curves for all models are presented in Figure 5.

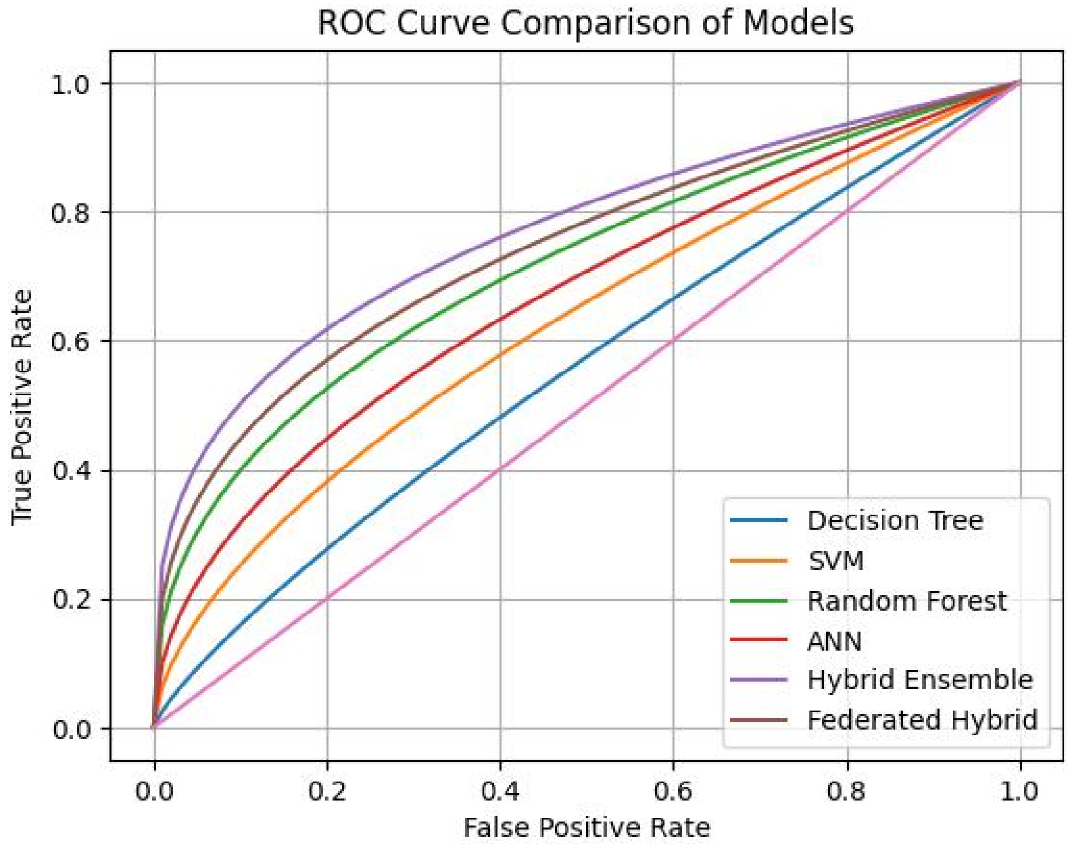

Figure 5: Receiver Operating Characteristic (ROC) curves comparing the classification performance of Decision Tree, SVM, Random Forest, ANN, the proposed model, and Federated the proposed models.

DISCUSSION

The findings of this study demonstrate that the proposed hybrid ensemble model significantly improves predictive performance in stroke telerehabilitation compared to individual machine learning models. Among the individual classifiers, Random Forest achieved the highest performance with an accuracy of 89.3% ± 1.2, followed by Artificial Neural Network and Support Vector Machine, while the Decision Tree model showed comparatively lower performance due to its sensitivity to data variability. These results are consistent with previous research indicating that ensemble-based models provide improved robustness and generalization in healthcare prediction tasks. The proposed hybrid ensemble model achieved the highest overall performance with an accuracy of 92.1% ± 1.0 and ROC-AUC of 0.93 ± 0.01. The improvement of approximately 2–3% over the best individual model confirms the effectiveness of combining complementary classifiers. The reduced standard deviation across validation folds indicates improved model stability and generalization, which are critical requirements for clinical applications. This improvement can be attributed to the ability of the ensemble model to balance bias and variance by integrating models with different learning characteristics.

The federated hybrid ensemble model achieved an accuracy of 90.8% ± 1.3 and ROC-AUC of 0.92 ± 0.01, demonstrating only a marginal performance reduction compared to the centralized model. This finding highlights the effectiveness of federated learning in preserving predictive performance while ensuring data privacy. The slight increase in variability observed in the federated model can be explained by data heterogeneity across distributed nodes, which reflects real-world healthcare scenarios. Despite this, the model maintained stable performance across cross-validation folds, confirming its robustness in decentralized environments. The convergence behavior of the federated learning model further supports its stability, with training loss decreasing consistently and stabilizing after approximately 20 communication rounds. This indicates efficient learning and reliable aggregation using the Federated Averaging algorithm. The ability to achieve near-centralized performance under distributed conditions is particularly important for healthcare systems where data sharing is restricted due to privacy and regulatory constraints. The confusion matrix analysis demonstrates balanced classification performance with high true positive and true negative rates, indicating reliable detection of rehabilitation outcomes. The model maintains strong sensitivity and specificity, reducing both false negatives and false positives. From a clinical perspective, minimizing false negatives is particularly important, as missed predictions may delay intervention and negatively impact patient recovery.

The ROC-AUC analysis further confirms the strong discriminative ability of the proposed model. The hybrid ensemble model achieved the highest ROC-AUC, indicating its effectiveness in distinguishing between outcome classes across different decision thresholds. The federated model showed comparable ROC performance, reinforcing the reliability of decentralized learning. The SHAP-based feature importance analysis provides additional validation of the model by identifying clinically relevant predictors such as motor function score and therapy adherence as the most influential features. Age and stroke severity also contributed significantly to predictions. The alignment of these findings with established clinical knowledge strengthens the interpretability and credibility of the model, which is essential for adoption in clinical practice. Overall, the results demonstrate that integrating hybrid machine learning, federated learning, and explainable AI provides a balanced solution that achieves high predictive accuracy, preserves data privacy, and ensures interpretability. This combination addresses key limitations of traditional machine learning approaches and supports the development of scalable and clinically applicable telerehabilitation systems, particularly in resource-constrained environments.

CLINICAL IMPLICATIONS

The proposed framework has important implications for clinical practice. It supports early prediction of rehabilitation outcomes, allowing timely intervention, and enables personalized therapy planning based on patient-specific data. The use of explainable AI provides interpretable outputs, which improves clinician confidence in model-driven decisions. In addition, the framework allows collaboration across institutions without sharing sensitive patient data, addressing key privacy concerns. It also improves access to rehabilitation services in resource-constrained environments. The integration of federated learning and explainable AI makes the system suitable for deployment in real-world clinical settings, particularly in developing countries.

LIMITATIONS

This study has several limitations that should be considered. First, although federated learning improves data privacy, it introduces challenges related to communication efficiency and synchronization across distributed nodes. Second, the study focuses on model-level performance and does not evaluate long-term clinical outcomes or patient-level improvements. Finally, the implementation of federated learning systems requires infrastructure and coordination across institutions, which may limit adoption in certain settings.

FUTURE DIRECTIONS

Future research should focus on validating the proposed framework using real-world clinical datasets collected from multiple healthcare institutions to improve generalizability and clinical applicability. Further work should explore the integration of wearable and sensor-based data to enable continuous patient monitoring and real-time feedback. In addition, the development of adaptive federated learning models capable of handling dynamic and heterogeneous data environments is needed to enhance system robustness. Future studies should also evaluate long-term rehabilitation outcomes and overall clinical effectiveness to establish practical impact. Addressing challenges related to interoperability and data standardization will be essential for system integration across different healthcare platforms. Moreover, policy-level considerations, including data governance, regulatory compliance, and ethical standards, should be examined to support large-scale deployment of such systems.

CONCLUSION

This study presents a hybrid machine learning and federated learning framework for stroke telerehabilitation that integrates explainable AI and privacy-preserving mechanisms. The findings show that hybrid ensemble models improve predictive performance, while federated learning maintains model effectiveness without compromising data privacy. In addition, explainable AI enhances interpretability and supports clinical usability. The proposed system offers a scalable and practical approach for stroke rehabilitation in resource-constrained environments, enabling the development of personalized and reliable telerehabilitation solutions through improved accuracy, privacy, and transparency.

REFERENCES

1. Dwork, Cynthia, and Aaron Roth. 2014. 'The algorithmic foundations of differential privacy', *Foundations and trends® in theoretical computer science*, 9: 211-487.
2. Feigin, Valery L, Michael Brainin, Bo Norrving, Sheila Martins, Ralph L Sacco, Werner Hacke, Marc Fisher, Jeyaraj Pandian, and Patrice Lindsay. 2022. 'World Stroke Organization (WSO): global stroke fact sheet 2022', *International journal of stroke*, 17: 18-29.
3. Kavakiotis, Ioannis, Olga Tsave, Athanasios Salifoglou, Nicos Maglaveras, Ioannis Vlahavas, and Ioanna Chouvarda. 2017. 'Machine learning and data mining methods in diabetes research', *Computational and structural biotechnology journal*, 15: 104-16.
4. Khurshid, Hamna, Urooj Khan, Kiran Ishaq, Amna Rafique, Mehr Un Nisa, Warda Khan, Kiran Hassan, Etisam Wahid, Shahzad Ahmad, and Ayesha Nisar. 2025. 'Effect of Task-Oriented Motor Training on Upper Limb Functional Recovery in Post-Stroke Patients: A Randomized Controlled Trial', *Link Medical Journal*: 1-13.
5. Lundberg, Scott M, and Su-In Lee. 2017. 'A unified approach to interpreting model predictions', *Advances in neural information processing systems*, 30.
6. McMahan, Brendan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. 2017. "Communication-efficient learning of deep networks from decentralized data." In *Artificial intelligence and statistics*, 1273-82. Pmlr.
7. Organization, World Health. 2024. *Rehabilitation 2030: meeting report, Geneva, Switzerland, 10-11 July 2023* (World Health Organization).
8. Rajkomar, Alvin, Jeffrey Dean, and Isaac Kohane. 2019. 'Machine learning in medicine', *New England Journal of Medicine*, 380: 1347-58.
9. Ribeiro, Marco Tulio, Sameer Singh, and Carlos Guestrin. 2016. "" Why should i trust you?" Explaining the predictions of any classifier." In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 1135-44.
10. Shah, Syed Awais, Nizar Ahmad, Waleed Muhammad, Lalina Zaib, Numan Khan,

Maaz Ali Mumtaz, Etisam Wahid, and Shahzad Ahmad. 2026. 'A HYBRID ENSEMBLE MACHINE LEARNING FRAMEWORK FOR EARLY PREDICTION OF TYPE 2 DIABETES: A STRUCTURED APPROACH TO PREDICTIVE HEALTHCARE', *Spectrum of Engineering Sciences*, 4: 1020-33.