Review of Computer Engineering Research
2026 Vol. 13, No. 4, pp. 196-219
ISSN(e): 2410-9142
ISSN(p): 2412-4281
DOI: 10.18488/76.v13i4.5353
© 2026 Conscientia Beam. All Rights Reserved.

FNST-DT: A federated neuro-symbolic framework for Parkinson's disease progression

Subhashree Balasubramanian[1]
Nallusamy Sellappa[2+]
Sobya Devaraj[3]
Manjuleeswari Pandy Sekaran[4]
Partha Sarathi Chakraborty[5]

[1,2,5]*School of Engineering Management and Continuing Education, Jadavpur University, Kolkata 700032, India.*
[1]*Email: msbalusubha@gmail.com*
[2]*Email: snallusamy@jadavpuruniversity.in*
[5]*Email: p_s_c2001@yahoo.com*
[3]*Department of Computer Science and Engineering, Institute of Engineering and Management, University of Engineering and Management, Kolkata 700091, India.*
[3]*Email: sobyadevaraj@gmail.com*
[4]*Department of Information Technology, Panimalar Engineering College, Chennai, Tamil Nadu 600123, India.*
[4]*Email: manjuleeswari09@gmail.com*


(+ Corresponding author)

Article History
Received: 24 June 2026
Revised: 7 October 2026
Accepted: 22 October 2026
Published: 9 November 2026

Keywords
Digital twin
Explainable artificial intelligence
Federated learning
Multimodal data analysis
Neuro-symbolic AI
Parkinson's disease
Privacy-preserving healthcare
Transformer networks.

ABSTRACT

Parkinson's disease (PD) progression modeling remains a challenging research problem due to variations in patient conditions, temporal changes in clinical characteristics, and the need for secure handling of healthcare data. Existing centralized deep learning approaches often face limitations related to data privacy, limited transparency, and difficulties in collaborative analysis across healthcare institutions. This study introduces FNST-DT, a Federated Neuro-Symbolic Framework with Digital Twin integration for privacy-preserving Parkinson's disease progression modeling. The framework integrates federated learning for decentralized model training, transformer-based temporal learning for extracting progression patterns from longitudinal clinical data, and neuro-symbolic reasoning for improving the interpretability of prediction outcomes. In addition, the digital twin module provides a patient-specific representation of disease progression to support personalized monitoring and clinical decision-making. The proposed framework was evaluated using the Parkinson's Telemonitoring dataset under a simulated federated learning environment. The results show that FNST-DT achieved a prediction accuracy of 94%, which is higher than MLP-LSVM (82%), PFMOO (85%), and PQFL (88%). The framework also obtained 95% privacy preservation and 93% computational efficiency, indicating its ability to balance predictive performance, security, and computational requirements. The findings suggest that FNST-DT can support future healthcare applications by enabling secure multi-institutional collaboration, continuous patient monitoring, and more transparent clinical decision support for Parkinson's disease management.

Contribution/Originality: This study contributes to the existing literature by integrating federated learning, transformer-based temporal modeling, neuro-symbolic reasoning, and digital twin technology into a unified framework for privacy-preserving Parkinson's disease progression modeling, enabling decentralized learning, interpretable prediction, personalized monitoring, and privacy-aware clinical decision support for healthcare applications.

1. INTRODUCTION

Parkinson's disease is a progressive and chronic neurodegenerative condition with motor symptoms such as tremors, rigidity, bradykinesia, postural instability, and non-motor symptoms like cognitive impairment, sleep

© 2026 Conscientia Beam. All Rights Reserved.

disturbances, and emotional disorders [1]. Patient-specific and heterogeneous progression patterns of PD complicate the accurate modeling of disease courses and the implementation of timely interventions [2]. As healthcare data collection from wearable sensors, telemonitoring systems, and clinical records grows rapidly, there is increasing potential to use AI to predict the progression of PD [3]. The classic centralized AI models have serious drawbacks related to data privacy, explainability, and scalability across multiple healthcare facilities [4]. The latest developments in deep learning, especially transformer-based models, have shown impressive performance in capturing long-range temporal dependencies in sequential data [5].

These models are highly applicable for analyzing longitudinal PD datasets, where time trends are important for understanding disease progression [6]. Data-driven methods are often not interpretable, and clinicians are not confident or able to interpret models' predictions in critical healthcare settings [7]. Neuro-symbolic AI has emerged through the combination of neural networks and symbolic reasoning to improve interpretability and incorporate domain knowledge [8]. Of particular interest, federated learning, as a privacy-preserving model, enables collaborative training of models on decentralized data without sharing raw patient data [9]. In healthcare environments, the safety and compliance of information are very critical [10].

The real-world object digital twin concept has opened up new possibilities in personalised medicine. It can simulate and trace the progress of a disease in a person in real time [11]. Management requires early diagnosis and constant monitoring. However, this is difficult to achieve because the onset is insidious and the manifestations are diverse [12]. The recommended architecture combines the benefits of the federated learning technique, the transformer-based temporal modelling method, and the neuro-symbolic reasoning method [13]. Digital twin technology can model the development of PD in a privacy-sensitive, explainable, and personalised way [14]. The FNST-DT framework aims to increase predictive accuracy and interpretability in the real-world healthcare context through a mixture of these techniques [15, 16].

Problem Statement: Parkinson's disease progression modeling remains a challenging task because of heterogeneous disease manifestations, complex temporal dependencies, multimodal clinical information, and strict privacy requirements associated with medical data. Although recent artificial intelligence models have improved prediction performance, many existing approaches rely on centralized learning architectures, provide limited clinical interpretability, and are not well suited for collaborative healthcare environments where secure data sharing is essential.

Research Gap: Recent studies have independently explored federated learning, transformer architectures, explainable artificial intelligence, and digital twin technologies for healthcare applications. However, these approaches are generally developed in isolation and do not simultaneously address privacy preservation, temporal disease progression, interpretable clinical reasoning, and patient-specific monitoring within a unified framework. Furthermore, limited attention has been given to integrating neuro-symbolic reasoning with federated transformer models to support transparent and personalized Parkinson's disease progression modeling.

Research Objective: Motivated by these limitations, this study aims to develop a Federated Neuro-Symbolic Framework with Digital Twin integration for privacy-preserving Parkinson's disease progression modeling. The proposed framework combines federated collaborative learning, transformer-based temporal feature extraction, neuro-symbolic reasoning, and patient-specific digital twin modeling to improve predictive performance, model interpretability, and personalized clinical decision support while maintaining data privacy.

This study contributes to the existing literature by introducing a federated neuro-symbolic framework for privacy-preserving Parkinson's disease progression modeling. This study uses a novel methodological integration of federated learning, transformer-based temporal learning, neuro-symbolic reasoning, and digital twin technology. This study is one of the very few investigations that combine these complementary paradigms within a unified clinical decision-support framework. The paper's primary contribution is demonstrating that this integrated framework

© 2026 Conscientia Beam. All Rights Reserved.

improves prediction accuracy, interpretability, privacy preservation, and personalized disease monitoring for intelligent healthcare applications.

The remainder of this paper is organized as follows. Section 2 discusses the related work and recent developments in Parkinson's disease progression modeling. Section 3 presents the proposed FNST-DT framework, including the federated learning architecture, transformer-based temporal modeling, neuro-symbolic reasoning, digital twin integration, and clinical decision support mechanism. Section 4 describes the experimental results and provides a detailed discussion of the obtained findings. Finally, Section 5 summarizes the conclusions and highlights the future research directions.

2. RELATED WORKS

Recent advances in artificial intelligence have significantly improved the prediction and monitoring of Parkinson's disease (PD) by enabling more accurate, privacy-aware, and interpretable healthcare solutions. Researchers have explored federated learning, transformer-based architectures, explainable artificial intelligence (XAI), multimodal deep learning, and privacy-preserving frameworks to overcome the limitations of conventional centralized learning systems. Although these approaches have demonstrated promising performance, most existing studies address only specific aspects of disease prediction, leaving several practical challenges unresolved. The following studies critically examine recent developments and identify the remaining research gaps that motivate the proposed framework.

Vaghasiya, et al. [17] propose a federated learning framework along with explainable AI for the diagnosis of Parkinson's disease using federated inventive MRI datasets. The proposed framework can train the model collaboratively and ensure that no confidential patient information is exchanged between institutions. XAI techniques will enhance the transparency of predictions and help clinicians with interpretation. While this approach has a lot of benefits, it basically deals with just diagnostic classification. It doesn't deal with the temporal progression of the disease. The lack of patient-specific monitoring further restricts its ability to manage chronic illnesses.

The prediction framework for differential privacy with ensemble machine learning models is proposed by Hussain, et al. [18]. This study applied extreme gradient boosting (XGB), random forest (RF), light gradient boosting machine (LGBM), and CatBoost (CAT) classifiers to clinical datasets, while Pareto frontier optimization was used to trade off predictive accuracy against privacy preservation. While the framework effectively enhances data privacy without a significant loss in prediction performance, its reliance on multiple ensemble models adds to the computational complexity and reduces scalability in large, distributed healthcare applications. Furthermore, the framework does not directly consider the evolution of the disease over time or clinically interpretable reasoning.

The Privacy-preserving Quantum Federated Learning (PQFL) framework was proposed for Parkinson's disease diagnosis by integrating federated learning, edge computing, and quantum-inspired encryption techniques, as described by Gajalakshmi and Kumar [19]. Utilizing phonation and motor activity data, the suggested approach allows local model training, along with the secure sharing of encrypted model parameters, as opposed to patient data. The framework efficiently lowers communication latency and reinforces privacy protection. The complexity associated with quantum-inspired encryption results in increased implementation overhead, while the framework is also primarily directed at disease diagnosis rather than progression modelling or personalized patient monitoring.

A study by Alshammari, et al. [20] proposed an autism spectrum disorder prediction framework based on federated learning that applies exploratory data analysis, machine learning models, federated optimization, and LIME. The inclusion of nonlinear learning models helps in achieving better prediction accuracy while protecting patient privacy through model training in a decentralized manner. Furthermore, explainability techniques improve the transparency of clinical decisions. Still, the proposed system was built specifically for autism screening; it does not consider the temporal characterisation, progression pattern, or patient-specific monitoring needs of Parkinson's disease. As a result, it has limited application for longitudinal prediction of neurodegenerative disease.

© 2026 Conscientia Beam. All Rights Reserved.

Akram and Chidananda [21] designed the Multimodal Adaptive Fusion Network (MAFNet) to enhance the prediction of Parkinson's disease staging and progression through the integration of clinical records, medical imaging, and genomic data. The approach incorporates temporal denoising, swarm intelligence-based feature selection, graph-attention multimodal fusion, and MLP-LSVM classification of multi-source data to obtain complementary disease characteristics. This multimodal strategy significantly improves prediction accuracy over using only one modality. However, relying on multiple data sources can complicate implementation and increase computational costs. Besides, because the framework adopts the centralized learning paradigm, it does not explicitly address data privacy, distributed model training, or collaborative learning across health institutions, all of which are critical for large-scale clinical deployment.

Ferdowsi, et al. [22] proposed a responsible artificial intelligence framework that emphasizes transparency, fairness, explainability, privacy preservation, and ethical compliance in healthcare prediction systems. The proposed architecture helps build AI solutions capable of clinical decision-making while complying with regulatory and ethical requirements. Even though the framework sets important design principles for responsible healthcare AI, its principal emphasis is on governance and ethics rather than on the predictive modeling of disease progression. Also, the framework does not include any learning mechanism with a temporal aspect or any personalized clinical reasoning.

Chen, et al. [23] developed a federated learning-based FedDBM framework for executing digital biomarker discovery across institutions. Through collaboration on model building, the proposed system does not require patient data to be centralized, thus improving compliance. By using explainability techniques, the interpretability of the predictions made by the algorithms is improved, and the identified biomarkers can be clinically validated. This framework is primarily focused on biomarker discovery rather than looking at disease progression over time. Furthermore, the system lacks the incorporation of patient-specific digital twin technology or neuro-symbolic reasoning, ultimately limiting its ability to generate personalized and continuously evolving clinical decision support.

Recent advances in transformer-based deep learning have demonstrated a remarkable capability in learning long-range temporal dependencies from sequential clinical data. Similarly, emerging healthcare foundation models have shown considerable potential for multimodal representation learning and large-scale medical knowledge integration. However, these approaches generally require substantial computational resources, large annotated datasets, and centralized training strategies, making their deployment challenging in privacy-sensitive distributed healthcare environments. In addition, most existing transformer and foundation model frameworks provide limited integration with symbolic clinical knowledge and personalized digital twin technologies. Consequently, there remains a need for an efficient framework that combines the predictive capability of transformer architectures with explainable reasoning, federated collaborative learning, and patient-specific disease modeling.

Ghazal, et al. [24] explored feature selection, multimodal data fusion, and optimization process assessment to enhance the diagnosis and progression analysis of Parkinson's disease. The results highlight that the integration of varied clinical information using machine learning helps improve the accuracy of diagnosis and personalized treatment. The research stresses the significance of transparent artificial intelligence and long-term clinical trials to enhance the generalizability of models. Even after all this, the proposed architecture relies on a centralised learning technique. Moreover, it provides privacy-preserving federated learning schemes that are not helpful for data sharing in healthcare scenarios.

Kumar, et al. [25] developed a diagnostic framework that is interpretable and privacy-preserving through the incorporation of SHAP-based explainability with isotonic regression to estimate confidence. The framework enhances prediction clarity and helps clinicians obtain more reliable explanations for diagnostic choices while concealing sensitive healthcare details. While the suggested method enhances model interpretability, it focuses on diagnostic functionality rather than explicitly modeling the progression of Parkinson's disease. Further, the ability to perform symbolic clinical reasoning and create personalized digital twins is not supported, which limits continuous disease monitoring.

© 2026 Conscientia Beam. All Rights Reserved.

Fritsch [26] developed a secure edge-based healthcare monitoring framework that employs wearable devices and home servers for continuous patient observation. The proposed architecture enables local processing of patient information while maintaining secure communication with healthcare providers through protected network channels. Experimental results demonstrated the feasibility of privacy-preserving remote healthcare monitoring with reduced communication overhead. However, the framework mainly addresses secure healthcare infrastructure and remote monitoring without incorporating advanced artificial intelligence techniques for disease progression prediction, personalized decision support, or explainable clinical reasoning.

Table 1. Comparison of prior studies.

Reference	Environment	Algorithm / Model	Privacy Preservation	Explainability	Personalization	Major Limitation
Vaghasiya, et al. [17]	Healthcare (MRI)	Federated Learning + Explainable AI	✓	✓	✗	Focused mainly on diagnosis without temporal disease progression.
Hussain, et al. [18]	Clinical Data	XGB, RF, LGBM, CAT + Differential Privacy	✓	✓	✗	Computational complexity limits scalability in large healthcare systems.
Gajalakshmi and Kumar [19]	Edge + Cloud	Quantum-Edge Federated Learning	✓	✗	✓	High implementation complexity for real-world deployment.
Alshammari, et al. [20]	Healthcare (ASD)	ML + Federated Learning + LIME	✓	✓	✗	Developed for ASD prediction rather than Parkinson's disease progression.
Akram and Chidananda [21]	Multimodal Clinical	MAFNet	✗	✗	✓	High computational cost and centralized learning architecture.
Ferdowsi, et al. [22]	AI Healthcare Systems	Responsible AI Framework	✓	✓	✗	Emphasizes governance rather than temporal disease prediction.
Chen, et al. [23]	Clinical + FL	FedDBM	✓	✓	✓	Limited to biomarker discovery without personalized monitoring.
Ghazal, et al. [24]	Multimodal Data	ML + Feature Optimization	✗	Partial	✓	Does not support privacy-preserving collaborative learning.
Kumar, et al. [25]	Healthcare	SHAP + Isotonic Regression	✓	✓	✗	Primarily designed for diagnostic interpretation.
Fritsch [26]	IoT + Edge	Smart Healthcare Monitoring	✓	Partial	✓	Monitoring framework without AI-based progression prediction.

© 2026 Conscientia Beam. All Rights Reserved.

Table 1 indicates that existing studies have made significant progress in privacy preservation, explainable artificial intelligence, multimodal learning, and distributed healthcare analytics. However, most approaches address these capabilities independently and therefore provide limited support for comprehensive Parkinson's disease progression modeling. Challenges related to temporal feature learning, integration of symbolic clinical knowledge, personalized patient monitoring, and unified privacy-preserving decision support remain largely unresolved. These observations provide the motivation for the proposed FNST-DT framework, which combines federated learning, transformer-based temporal modeling, neuro-symbolic reasoning, and digital twin technology within a single architecture.

3. MATERIALS AND METHODS

The progression of PD is hard to model because it is temporal, multimodal, and very patient-specific in nature. To guarantee privacy, interpretability, and personalization, the suggested FNST-DT system will utilize a federated learning algorithm, transformer-based sequence modeling, neuro-symbolic reasoning, digital twin technology, and intelligent decision support. It is an integrated architecture that can support decentralized learning among institutions, transparent clinical inference, real-time patient monitoring, and adaptive treatment recommendations for scalable and explainable management of Parkinson's disease.

3.1. Federated Neuro-Symbolic Learning Architecture that Preserves Privacy

The proposed FNST-DT framework integrates federated learning with neuro-symbolic reasoning within the local client-training pipeline. It enables decentralized model training across multiple virtual federated clients while preserving sensitive patient information through differential privacy and secure aggregation. Transformer-based temporal modeling with rule-based clinical reasoning improves prediction accuracy and interpretability, both of which are essential for overcoming the major drawbacks of traditional black-box deep learning systems in the analysis of Parkinson's disease.

The proposed FNST-DT framework introduces a decentralized and privacy-preserving approach for Parkinson's disease progression modeling in a federated learning environment. The publicly available Parkinson's Telemonitoring dataset is partitioned into multiple virtual federated clients to simulate decentralized healthcare institutions. Each virtual client processes longitudinal telemonitoring data and clinical information, including UPDRS scores and demographic attributes, through local preprocessing before transmitting only model updates for federated aggregation. The preprocessed data are subsequently provided as input to the core FNST-DT framework for disease progression modeling. The framework incorporates a transformer to learn long-range temporal dependencies and a neuro-symbolic reasoning module to fuse deep learning representations with rule-based clinical knowledge that is interpretable. A personalized digital twin is created to replicate real-time disease dynamics and allow mapping of its state at the patient level. Each institution trains and optimizes local models and transmits only encrypted model weights to the central federated aggregation server. The server performs global model aggregation using methods such as FedAvg, which creates one global model without violating data privacy. The explainable decision-making model is based on the aggregated model and offers interpretable predictions, transparent rule-based explanations, and digital twin-driven insights for treatment. This architecture, shown in Figure 1, ensures scalability, improves prediction accuracy, and enables intelligent, personalized, and safe clinical decision support in real-world healthcare systems.

© 2026 Conscientia Beam. All Rights Reserved.

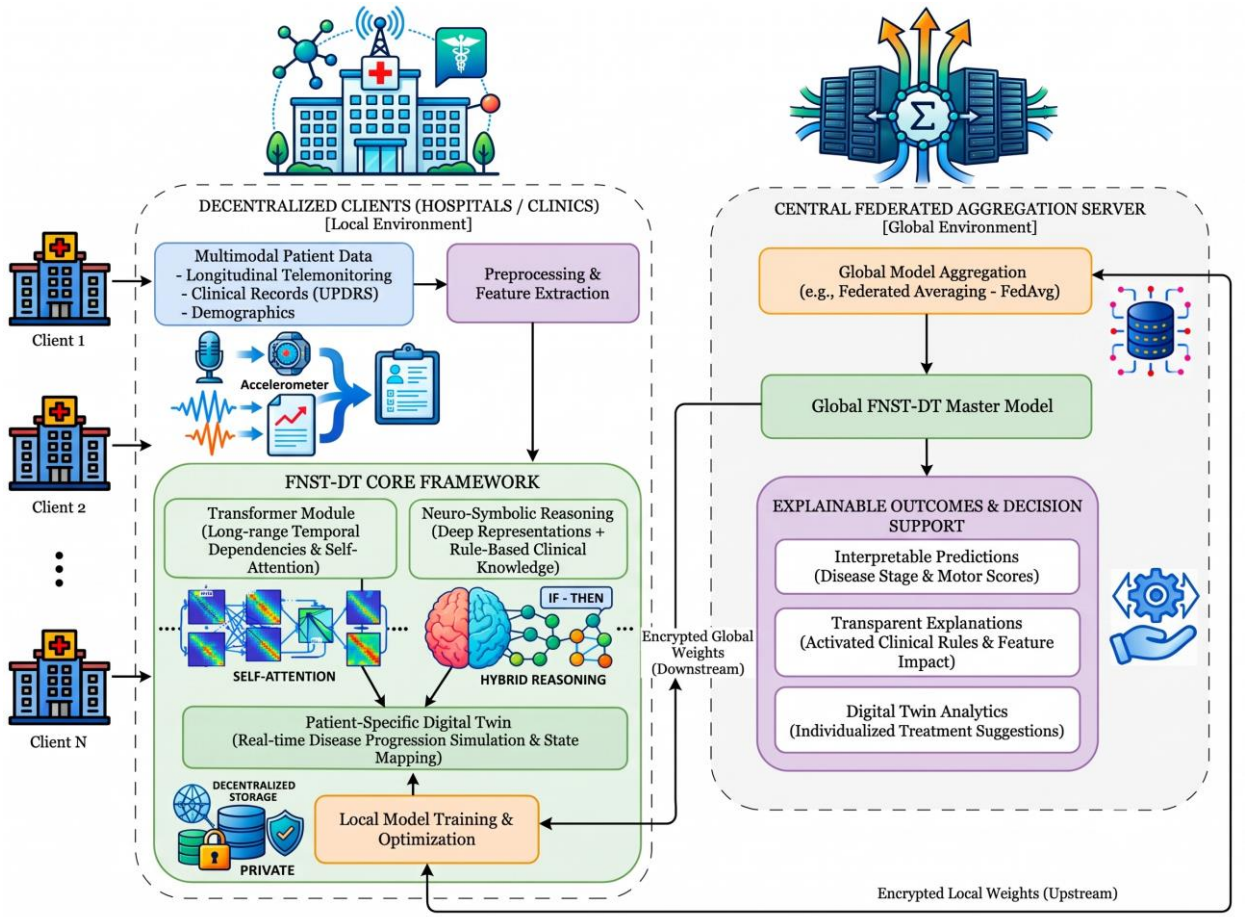


Figure 1. Architecture of federated neuro-symbolic transformer with digital twin integration.

Multimodal patient data representation X_i defined using Equation 1.

$$X_i = W_t \cdot T_i + W_c \cdot C_i + W_d \cdot D_i + b_x \quad (1)$$

The multimodal patient data vector X_i for patient i integrates heterogeneous healthcare inputs. The vector T_i represents longitudinal telemonitoring signals, C_i denotes clinical scores such as UPDRS, and D_i includes demographic features. The matrices, W_t, W_c, W_d control the contribution of each modality, while b_x is a bias term ensuring consistent feature scaling.

Local model training & optimization $\theta_i^{(t+1)}$ processed by Equation 2.

$$\theta_i^{(t+1)} = \theta_i^{(t)} - \eta \cdot \nabla \mathcal{L}_i\left(1 - \theta_i^{(t)}\right) \quad (2)$$

The local model parameters θ_i for institution i are updated iteratively using gradient descent. The learning rate η controls the update step size, and $\nabla \mathcal{L}_i$ represents the gradient of the local loss function based on patient data. This formulation enables each virtual federated client to independently optimize its local model without sharing raw patient data.

Global model aggregation: federated averaging $\theta^{(t+1)}$ evaluated using Equation 3.

$$\theta^{(t+1)} = \sum_{i=1}^{K} \frac{n_i}{\sum_{j=1}^{K} n_j} \cdot \theta_i^{(t)} \quad (3)$$

© 2026 Conscientia Beam. All Rights Reserved.

The model parameters $\theta^{(t+1)}$ for the global model are acquired by pooling the local model updates of K institutions. In institution i, the term n_i denotes the number of samples for weighted averaging. The FedAvg approach combines knowledge and preserves privacy.

Global FNST-DT master model is processed using Equation 4.

$$Y_t = \sigma(\text{Transformer}(X_t; \theta) + \mathcal{R}(X_t)) \quad (4)$$

The FNST-DT master model generates the output prediction Y_t. The transformer component uses global parameters θ to model long-range temporal dependencies, while $\mathcal{R}(X_t)$ captures neuro-symbolic reasoning that incorporates clinical rules. The activation function σ yields interpretable rankings of disease progression.

Patient-specific digital twin Z_i^{t+1} valued using Equation 5.

$$Z_i^{t+1} = Z_i^t + \alpha \cdot (f_\theta(X_i^t) - Z_i^t) \quad (5)$$

The digital twin state Z_i^{t+1} depicts the changing status of patient i. The notation f_θ stands for the prediction from a trained FNST-DT model, and Z_i^t is the twin state at time t. The parameter α affects the system's update rate, enabling smooth adaptation to new patient data. Furthermore, this enables real-time disease tracking.

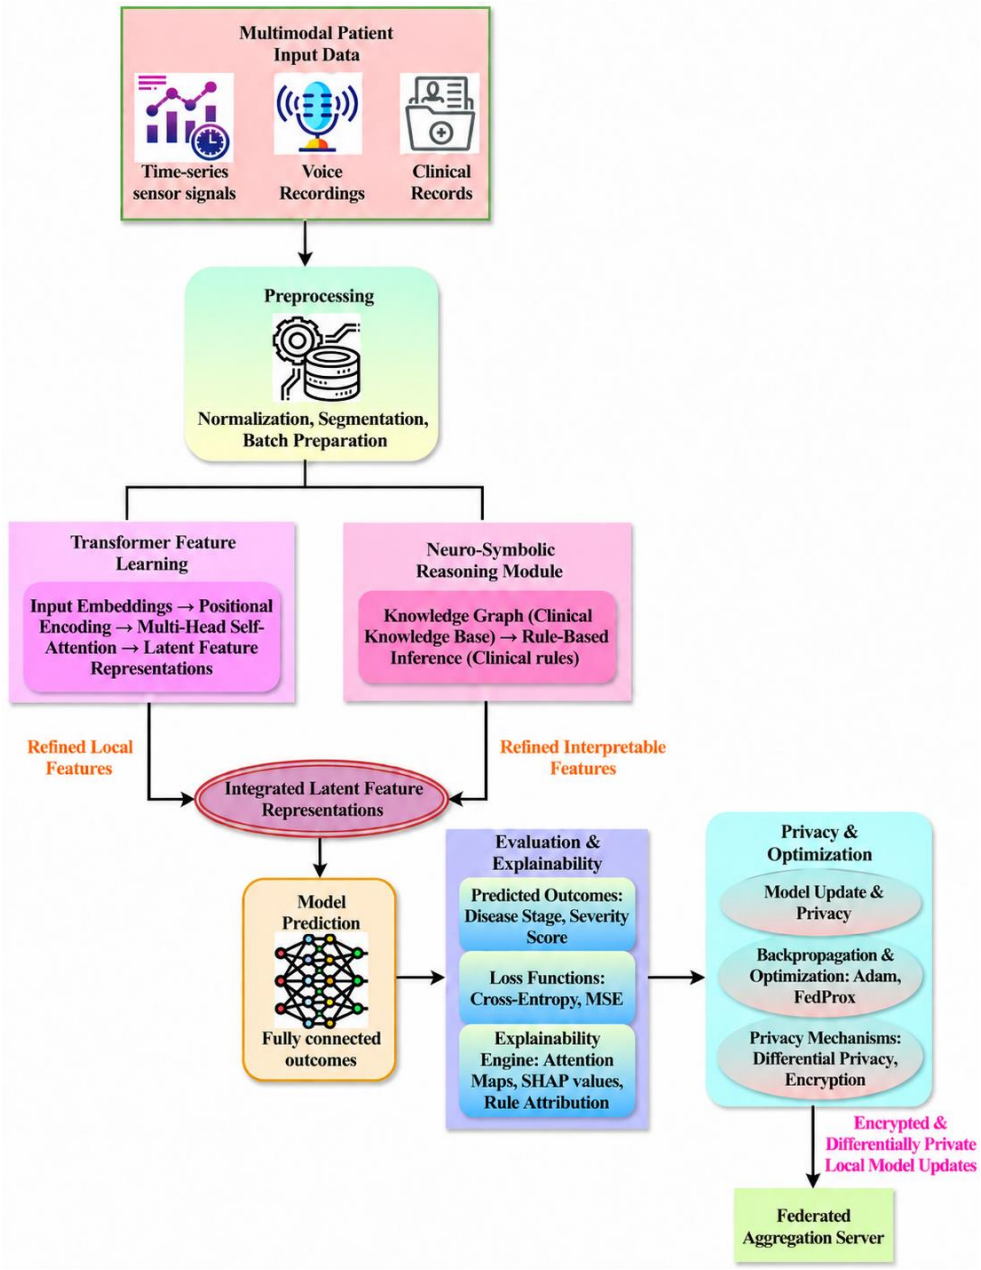


Figure 2. Continuous learning and adaptive model update framework in FNST-DT.

© 2026 Conscientia Beam. All Rights Reserved.

3.2. Digital Twin of the Patient to Model Diseases in Real Time

The proposed digital twin module provides a dynamic, patient-specific virtualization of Parkinson's disease progression. The combination of real-time and predictive strategies enables simulation of disease progression and potential treatments. It facilitates effective, personalized healthcare by enabling continuous monitoring, early action, and the evaluation of hypothetical scenarios, and aids clinical decision-making in challenging neurological disorders.

The FNST-DT framework's Local Client Training module is charged with patient-specific pattern learning and guarantees data privacy in each healthcare institution. It starts with multimodal patient data input, comprising time-series sensor signals, voice recordings, and clinical records. This data is normalized, segmented, and prepared in batches to achieve uniformity and effectiveness in processing. This is followed by a transformer-based feature learning component, in which input embeddings and positional encoding are employed to learn temporal dependencies. The multi-head self-attention model allows the model to learn long-range relationships in disease progression through the generation of latent feature representations. These characteristics are further refined with the help of a neuro-symbolic reasoning module that combines clinical knowledge bases and rule-based inference via a knowledge graph and generates interpretable and clinically relevant representations. The refined features are fed into fully connected outcomes to provide predictions such as disease stage or disease-severity score. The optimization is guided by loss functions, such as cross-entropy and mean squared error. An explainability engine provides transparency in the form of attention maps, SHAP values, and rule attribution. Lastly, it optimizes the model and backpropagates using algorithms such as Adam or FedProx. Figure 2 demonstrates that privacy mechanisms such as differential privacy and encryption can be implemented before transferring the updated model parameters to the federated server securely.

Multimodal patient input data X_i^t evaluated using Equation 6.

$$\mathrm{X}_i^t = \mathrm{W}_s \cdot \mathrm{S}_i^t + \mathrm{W}_v \cdot \mathrm{V}_i^t + \mathrm{W}_c \cdot \mathrm{C}_i^t + \mathrm{b}_x \tag{6}$$

The vector S_i^t represents time-series sensor signals, V_i^t corresponds to voice recordings, and C_i^t denotes clinical records such as diagnostic scores. The matrices $\mathrm{W}_s, \mathrm{W}_v, \mathrm{W}_c$ assign weights to each modality, while b_x ensures proper scaling.

Preprocessing normalization & segmentation $\mathrm{X}_i^{\mathrm{prep}}$ computed using Equation 7.

$$\mathrm{X}_i^{\mathrm{prep}} = \mathrm{Batch}\left(\mathrm{Norm}\left(\mathrm{X}_i^t - \mu_i\right)\right) \tag{7}$$

The term μ_i represents the mean used for centering patient data, while Norm ensures consistent scaling across features. The function Batch segments the data into mini-batches, enabling efficient and stable training of deep models.

Transformer feature learning H_t evaluated using Equation 8.

$$\mathrm{H}_t = \mathrm{softmax}\left(\frac{\mathrm{Q}_t \mathrm{K}_t^{\mathrm{T}}}{\sqrt{d_k}}\right)\mathrm{V}_t + \mathrm{P}_t \tag{8}$$

The matrices Q_t, K_t, and V_t denote query, key, and value representations derived from embeddings, while d_k is the scaling factor for stability. The term P_t represents positional encoding, which preserves sequence order. This formulation enables the model to learn long-range relationships in disease progression.

Integrated latent feature Z_t representation processed using Equation 9.

$$\mathrm{Z}_t = \sigma(\mathrm{W}_h \cdot \mathrm{H}_t + \mathrm{W}_r \cdot \mathcal{K}(\mathrm{H}_t) + \mathrm{b}_z) \tag{9}$$

The term H_t represents transformer-derived features, while $\mathcal{K}(\mathrm{H}_t)$ denotes outputs from the knowledge graph and rule-based inference module. The matrices W_h and W_r balance learned and symbolic contributions, and b_z is a bias term. The activation function σ ensures nonlinear integration, producing clinically meaningful representations.

The weighting matrices used in Equation 9 regulate the relative contributions of transformer-derived representations and neuro-symbolic reasoning during feature integration. During practical implementation, these parameters should be selected through validation-based hyperparameter tuning to achieve an appropriate balance between predictive accuracy, interpretability, and model stability.

© 2026 Conscientia Beam. All Rights Reserved.

Federated aggregation server secure model update $\theta^{(t+1)}$ processed using Equation 10.

$$\theta^{(t+1)} = \sum_{i=1}^{K} \frac{n_i}{\sum_{j=1}^{K} n_j} \cdot \left(\theta_i^{(t)} + \mathcal{N}(0, \sigma^2)\right) \tag{10}$$

The global model parameters $\theta^{(t+1)}$ are updated at the federated server by aggregating encrypted local models. The term $\theta_i^{(t)}$ represents parameters from client i, and n_i is the number of samples at that client. The noise term $\mathcal{N}(0, \sigma^2)$ introduces differential privacy, protecting sensitive information. This ensures secure and privacy-preserving global learning.

Algorithm 1: Federated Neuro-Symbolic Transformer Training

1. $Input: \{Di = (Xi, Yi)\}, \theta 0, \eta, R, H$
2. $Output: \theta_{global}$
3. $Initialize\ \theta_{global} \leftarrow \theta 0$
4. $for\ r = 1\ to\ R\ do$
5. $\quad for\ each\ client\ i \in \{1, \ldots, N\}\ do$
6. $\quad Zi \leftarrow Normalize(Xi) * We + P$
7. $\quad for\ h = 1\ to\ H\ do$
8. $\quad Qh \leftarrow Zi\, Wh^Q;\ Kh \leftarrow Zi\, Wh^K;\ Vh \leftarrow Zi\, Wh^V$
9. $\quad Ah \leftarrow Softmax((Qh\ Kh\hat{}T)/\sqrt{dk})$
10. $\quad Oh \leftarrow Ah\, Vh$
11. $\quad end\ for$
12. $\quad O \leftarrow Concat(O1, \ldots, OH)\, Wo$
13. $\quad Z'i \leftarrow ReLU(O\, W1 + b1)\, W2 + b2$
14. $\quad Ri \leftarrow Extract_{Rules(Xi)}$
15. $\quad Fi \leftarrow \alpha\, Z'i + (1 - \alpha)\, Ri$
16. $\quad \hat{Y}i \leftarrow Softmax(Fi\, Wc + bc)$
17. $\quad Li \leftarrow \Sigma\, Loss(\hat{Y}i, Yi)$
18. $\quad \theta i \leftarrow \theta_{global} - \eta\, \nabla Li$
19. $\quad Send\ \theta i\ to\ server$
20. $end\ for$
21. $\theta_{global} \leftarrow \Sigma i \left(\frac{ni}{\Sigma j\ nj}\right) \theta i$
22. $Broadcast\ \theta_{global}$
23. $end\ for$

The FNST-T Algorithm 1 is a decentralized training process for a transformer-based model combined with neuro-symbolic reasoning across multiple healthcare clients. First, a universal model is sent to all involved institutions. Each client preprocesses its local multimodal data and produces positionally encoded embeddings. The transformer module uses multi-head self-attention to learn long-range temporal dependencies, and feature refinement is then performed by a feed-forward network. A neuro-symbolic module derives clinical rules and fuses them with deep features to boost interpretability. The model subsequently generates predictions, and a loss function is produced and optimized. Using gradient descent, each client updates its local model parameters. These updates are securely sent to a central server, where federated aggregation with weighted averaging is performed based on the size of the datasets. The new world model is reallocated to all the clients in the second version. This recursive algorithm guarantees privacy, maintains scalability, and enhances predictive accuracy while maintaining interpretability across healthcare systems distributed over networks.

3.3. Smart and Understandable Clinical Decision Support System

The proposed decision support system incorporates both standard clinical treatment recommendations and personalized treatment recommendations. Furthermore, this is achieved through multimodal predictive models using symbolic inference. It carries out risk stratification, categorization of disease stages, and alert generation, with

© 2026 Conscientia Beam. All Rights Reserved.

transparency through explainable AI processes, including attention visualization and rule attribution. An increasingly involved human-in-the-loop ensures clinical validation, offering trust, flexibility, and functionality.

The IDSS of the FNST-DT model is meant to provide personalized, explainable, and clinically sound management recommendations for Parkinson's disease. The global federated model's outcomes, patient-specific simulations, and neuro-symbolic inference results are combined to produce a complete patient context. A data integration module integrates multimodal data, for example, historical and real-time trends in disease progression. A clinical knowledge processing layer is built into the system, which uses medical guidelines and rule-based engines to provide evidence-based decision-making. Risk stratification, disease stage classification, and the analysis of the progression trend are carried out by a prediction interpretation layer. The treatment recommendation engine creates individualized interventions, such as medication changes, treatment plans, and lifestyle recommendations, based on these insights. A warning mechanism is used to detect emergency situations and send alerts. The explainability module also offers easy-to-understand information with rule-based and attention-driven explanations. The human-in-the-loop interface enables clinicians to confirm the recommendations and ensure reliability, adaptability, and continuous improvement in clinical decision-making, as shown in Figure 3.

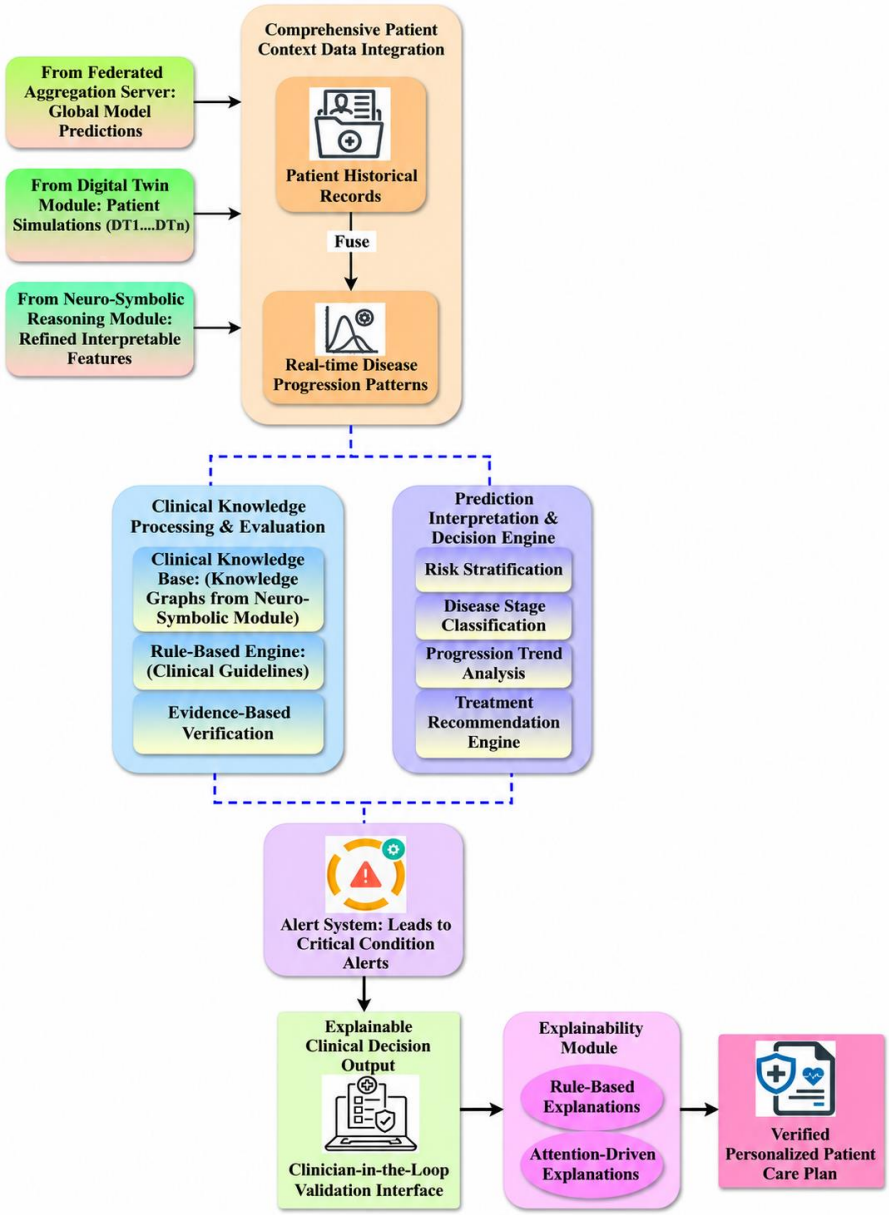


Figure 3. Intelligent decision support system in the FNST-DT framework.

© 2026 Conscientia Beam. All Rights Reserved.

Algorithm 2: Digital Twin Intelligent Decision Support System

1. $Input: Xi, \theta_{global}, Si_{DT}, Ri$
2. $Output: \hat{Y}i, Ti, Si_{DT}, new$
3. $Zi \leftarrow Xi\,We + P$
4. $for\ h = 1\ to\ H\ do$
5. $Qh \leftarrow Zi\,Wh^{Q};\ Kh \leftarrow Zi\,Wh^{K};\ Vh \leftarrow Zi\,Wh^{V}$
6. $Ah \leftarrow Softmax\left(\frac{Qh\,Kh^{T}}{\sqrt{dk}}\right)$
7. $Oh \leftarrow Ah\,Vh$
8. $end\ for$
9. $O \leftarrow Concat(O1, \ldots, OH)\,Wo$
10. $Z'i \leftarrow ReLU(O\,W1 + b1)\,W2 + b2$
11. $Ri \leftarrow Apply_{Clinical_{Rules(Xi)}}$
12. $Fi \leftarrow \alpha\,Z'i + (1-\alpha)\,Ri$
13. $\hat{Y}i \leftarrow Softmax(Fi\,Wc + bc)$
14. $if\ \max(\hat{Y}i) > \tau_{high}\ then$
15. $Risk \leftarrow High$
16. $else\ if\ \max(\hat{Y}i) > \tau_{mid}\ then$
17. $Risk \leftarrow Moderate$
18. $else$
19. $Risk \leftarrow Low$
20. $end\ if$
21. $Si_{DT}, new \leftarrow f(Si_{DT}, \hat{Y}i, Xi)$
22. $Ti \leftarrow Treatment(Risk, Si_{DT}, new)$
23. $return\ \hat{Y}i, Ti, Si_{DT}, new$

Transformer-based inference, neuro-symbolic reasoning, and digital twin simulation are combined in the DT-IDSS Algorithm 2 to offer personalized clinical decision-making support. It starts with patient-specific input data, which are converted into embeddings and fed into a multi-head attention mechanism to reflect temporal relationships. A feed-forward network is then used to refine the resulting features, and a neuro-symbolic fusion method is used in combination with clinical rules. Examples of predictions made by the model include the stage or severity level of the disease. These predictions are then used to define the patient's risk category according to predetermined thresholds. The computer representation of the patient, called a digital twin model, reflects the patient's current state and integrates current measurements with forecasted results to enable real-time disease tracking. Depending on the assessed level of risk and the knowledge offered by a digital twin, treatment proposals are tailored through medication and lifestyle changes. This ensures that clinicians obtain patient-specific, interpretable, data-driven guidance that is adaptable to the disease context while enabling proactive healthcare.

The FNST-DT framework's continuous learning and model update module ensures that the system is adaptive, accurate, and robust in dynamic clinical settings. It constantly gathers new patient data, updates to digital twins, prediction errors, and feedback provided by clinicians to train the model over time. A concept shift and data drift detection system follows the statistical variation in the incoming data distributions with the help of KL divergence and the population stability index. A retraining trigger is activated when a large drift is detected. Experience replay and buffering mechanisms are used to validate the incoming data, preprocess it, and combine it with historical data to avoid catastrophic forgetting. Fine-tuning and online updates are performed by an incremental learning engine based on gradient-based optimization, and model stability is ensured by regularization methods. The system also allows federated retraining, where an array of institutions can update the global model together without exchanging sensitive information. Performance evaluation metrics, such as accuracy, privacy preservation levels, and computational efficiency, are continually observed to evaluate model improvements. Clinician feedback will be used to increase model reliability and trust. The new model is implemented on edge devices and aligned with the digital twin, allowing real-time, adaptive, and personalized management of Parkinson's disease, as shown in Figure 4.

© 2026 Conscientia Beam. All Rights Reserved.

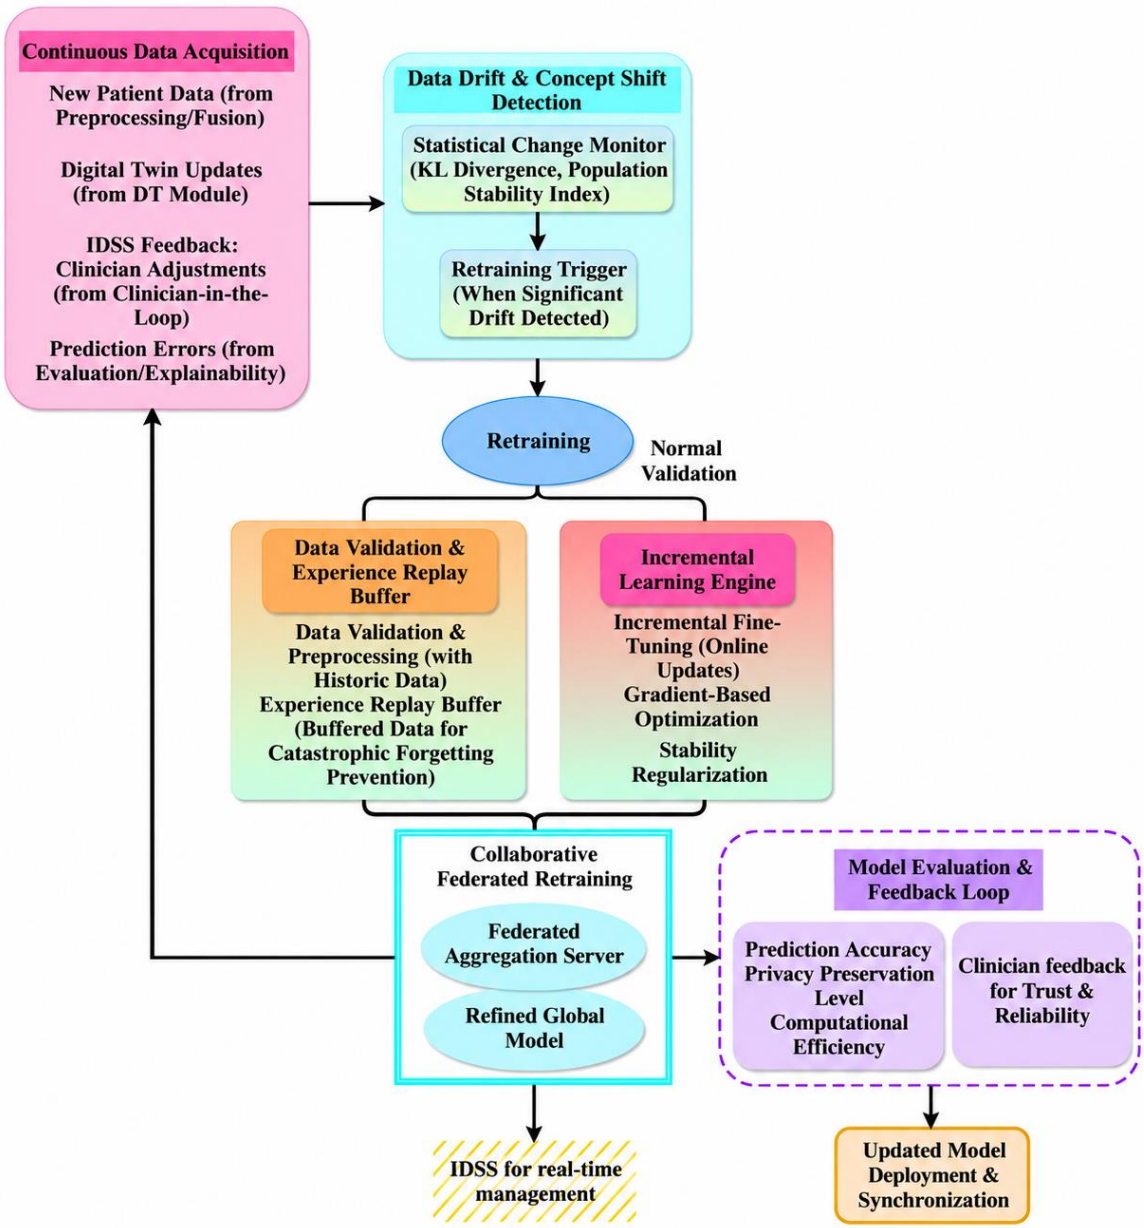


Figure 4. Continuous learning and model update module in the FNST-DT.

FNST-DT is an integrated platform that incorporates privacy-preserving federated learning, interpretable neuro-symbolic AI, and patient-specific digital twins for Parkinson's disease modeling. This technology provides sophisticated forecasts, real-time observations, and customized treatment suggestions with the help of a sophisticated decision support system. The design framework has extension potential, great power, and clinical interpretability; therefore, it can be transferred into practice to enable early diagnosis, sustained monitoring, and adaptive treatment in a distributed health care system.

3.4. Implementation Considerations

The proposed FNST-DT framework is presented as a conceptual architecture for privacy-preserving Parkinson's disease progression modeling. During practical deployment, implementation parameters, including the optimization algorithm, learning rate, batch size, transformer architecture, and other training hyperparameters, should be selected according to the characteristics of the target clinical dataset and available computational resources.

The optimization process should be guided by validation-based hyperparameter tuning to identify an appropriate configuration that balances predictive accuracy, convergence stability, computational efficiency, and privacy

© 2026 Conscientia Beam. All Rights Reserved.

preservation. Similarly, the weighting coefficients employed in the mathematical formulations should be selected through empirical validation to ensure reliable model performance and robust convergence under different deployment scenarios.

4. RESULTS AND DISCUSSION

Dataset Description: The Unlocking Clues to Parkinson's Disease Progression dataset (on Kaggle) is a rich telemonitoring dataset aimed at predictive modeling of Parkinson's disease using longitudinal clinical data. It includes important characteristics, such as motor and total UPDRS scores. These are commonly used to measure the severity and progression of the disease. The dataset enables researchers to examine temporal patterns, build machine learning models, and improve early diagnosis and tracking methods for Parkinson's disease progression [16]. Table 2 explains the simulation environment of the given dataset.

Before model development, duplicate records were removed, and the dataset was examined for missing values. Missing entries were handled using median-based imputation, followed by min-max normalization to ensure numerical consistency among clinical features. The dataset was randomly divided into 80% for training and 20% for testing, while 10% of the training set was reserved for validation during model optimization. To improve reproducibility, all experiments were conducted using a fixed random seed of 42.

Table 2. The simulation environment.

Parameter	Description
Dataset	Parkinson's Telemonitoring Dataset (Kaggle)
Number of Samples	~5,800+ records
Features	Biomedical voice measurements, UPDRS scores, temporal attributes
Data Type	Multimodal, time-series clinical data
Preprocessing Techniques	Duplicate removal, median imputation, Min-Max normalization, feature scaling
Framework	FNST-DT (Federated Neuro-Symbolic Transformer with Digital Twin)
Programming Language	Python
Libraries/Tools	TensorFlow / PyTorch, Scikit-learn, NumPy, Pandas
Federated Learning Setup	Five virtual federated clients simulated from the public Parkinson's Telemonitoring dataset
Communication Protocol	Secure aggregation (e.g., FedAvg)
Transformer Configuration	Multi-head attention, positional encoding, sequence modeling
Neuro-Symbolic Module	Rule-based clinical reasoning integrated with deep learning
Digital Twin Environment	Patient-specific virtual model with real-time updates
Hardware	Intel i7 Processor / GPU (NVIDIA RTX series)
Train/Test Split	80% Training / 20% Testing
Validation Strategy	10% of training data used for validation
Random Seed	42
Federated Client Partition	Five virtual clients created from the public dataset

The analysis of the FNST-DT framework is based on the main performance indicators, such as prediction accuracy, privacy preservation, computational efficiency, training convergence, and communication overhead. These metrics are evaluated over different training epochs, feature sizes, client distributions, and communication rounds. To confirm the effectiveness of the proposed model in delivering accurate, privacy-preserving, and scalable Parkinson's disease progression modeling in distributed healthcare settings, the proposed model is compared to the current models (MLP-LSVM, PFMOO, and PQFL).

© 2026 Conscientia Beam. All Rights Reserved.

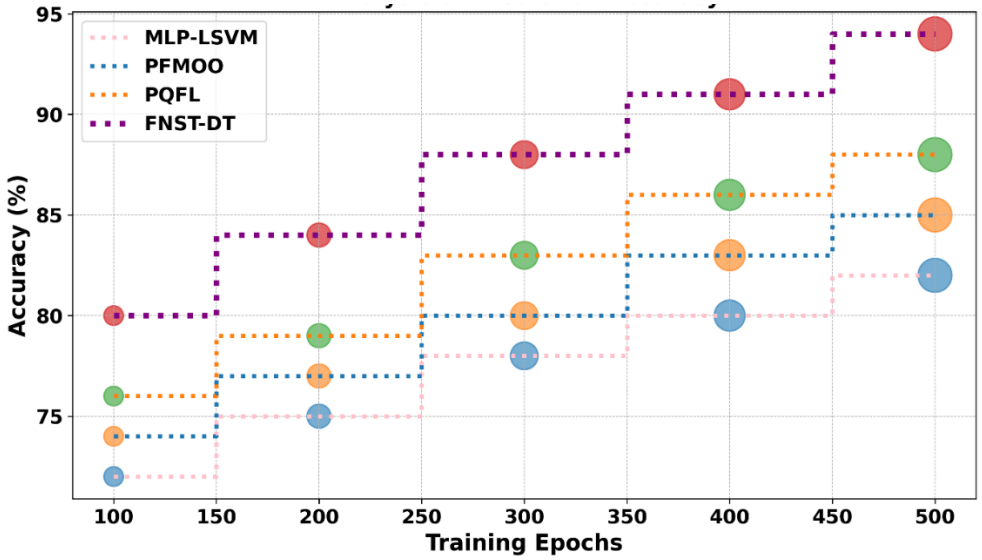


Figure 5. Analysis of prediction accuracy.

4.1. Prediction Accuracy

The comparison of prediction accuracy over the training epochs (100–500) shows that all models steadily improve, with the proposed FNST-DT demonstrating the best performance. MLP-LSVM, PFMOO, PQFL, and FNST-DT have accuracies of 72, 74, 76, and 80, respectively, at 100 epochs. Accuracies grow to 78, 80, 83, and 88 as training reaches 300 epochs. FNST-DT achieves the best accuracy of 94 at 500 epochs, compared with PQFL (88%), PFMOO (85%), and MLP-LSVM (82%). The gradual increase shows that the models learned the temporal patterns successfully, whereas the steeper increase in FNST-DT demonstrates that it learned the complicated dependencies. Such findings affirm that the suggested framework highly improves predictive performance over current approaches, as presented in Figure 5.

Analysis of prediction accuracy PA valued using Equation 11.

$$\mathrm{PA} = \frac{1}{N}\sum_{i=1}^{N} \mathbb{I} * \hat{y}_i - y_i \qquad (11)$$

Here, N denotes the total number of instances, y_i is the true label, and $\hat{y}_i$ is the predicted label for instance i. The indicator function $\mathbb{I}$ returns 1 when the prediction matches the ground truth and 0 otherwise, providing a direct measure of model correctness.

The improved prediction capability of FNST-DT is consistent with recent studies that reported the effectiveness of transformer-based models in capturing temporal patterns from Parkinson's disease data. However, many existing approaches mainly focus on centralized prediction and do not address privacy-preserving collaborative learning. In contrast, FNST-DT integrates transformer-based temporal learning with federated learning and neuro-symbolic reasoning to achieve accurate and interpretable disease progression modelling [27, 28].

Table 3. Comparison of model robustness.

Noise Level (%)	MLP-LSVM (%)	PFMOO (%)	PQFL (%)	FNST-DT (Proposed) (%)
10	79.3 ± 2.0	82.7 ± 1.8	86.1 ± 1.6	90.8 ± 1.3
20	75.6 ± 2.3	79.1 ± 2.0	83.4 ± 1.8	88.2 ± 1.5
30	71.2 ± 2.6	75.8 ± 2.2	80.5 ± 2.0	85.6 ± 1.7
40	66.8 ± 2.9	72.3 ± 2.5	77.6 ± 2.3	82.9 ± 1.9
50	62.5 ± 3.2	68.7 ± 2.8	74.2 ± 2.6	79.5 ± 2.1

© 2026 Conscientia Beam. All Rights Reserved.

4.2. Model Robustness

Performance deterioration is observed across all models as noise increases, as shown in a robustness analysis of different noise levels (10–50%), but FNST-DT is more stable. When the noise level is 10%, the accuracies achieved by MLP-LSVM, PFMOO, PQFL, and FNST-DT are 79.3%, 82.7%, 86.1%, and 90.8%, respectively. At a noise level of 30%, performance decreases to 71.2%, 75.8%, 80.5%, and 85.6%. FNST-DT attains 79.5% at the maximum noise level of 50%, surpassing the competing methods PQFL (74.2%), PFMOO (68.7%), and MLP-LSVM (62.5%). In addition, FNST-DT has lower deviation values, making it more consistent and reliable. The findings reveal that the proposed model can effectively deal with noisy and uncertain data due to its federated architecture and transformer-based learning. In general, FNST-DT is highly robust and can be used in real-life healthcare settings where data are imperfect, as shown in Table 3.

Comparison of model robustness MR computed using Equation 12.

$$\mathrm{MR} = 1 - \frac{1}{N}\sum_{i=1}^{N} \frac{\|f(\mathrm{x}_i) - f(\mathrm{x}_i + \delta_i)\|_2}{\|f(\mathrm{x}_i)\|_2 + \epsilon} \tag{12}$$

The function f represents the model output; x_i is the original input, and δ_i is a small perturbation or noise. The normalized difference measures how much predictions change under disturbance. A higher MR indicates more stable and robust behavior.

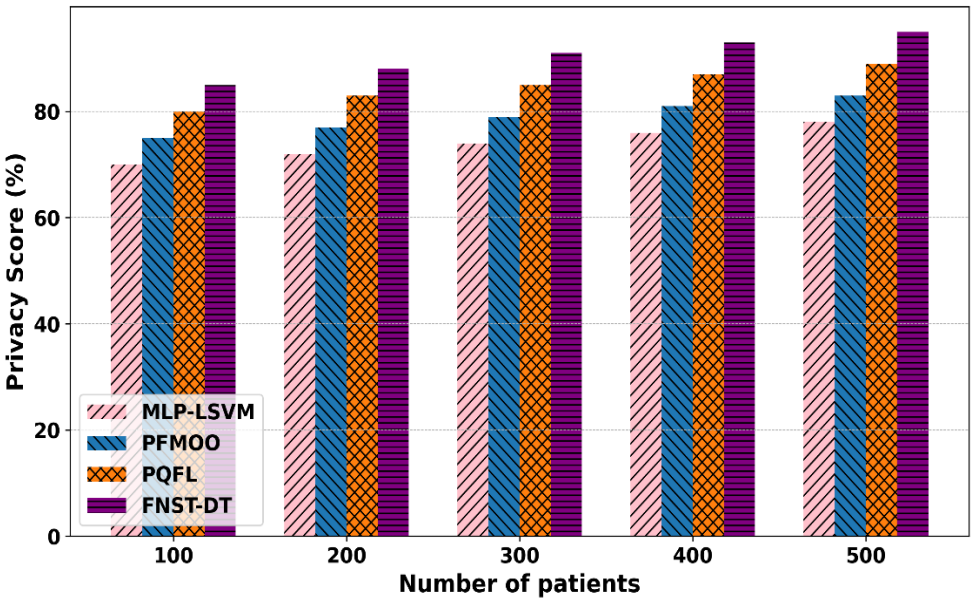


Figure 6. Analysis of privacy preservation level.

4.3. Privacy Preservation Level

The comparison of the level of privacy preservation for different numbers of clients (100–500 clients) indicates a definite increase in all models, with the highest privacy scores being achieved by FNST-DT. At 100 clients, MLP-LSVM, PFMOO, PQFL, and FNST-DT record 70%, 75%, 80%, and 85%, respectively. As the number of clients increases to 300, the scores improve to 74%, 79%, 85%, and 91%. At 500 clients, FNST-DT reaches 95%, outperforming PQFL (89%), PFMOO (83%), and MLP-LSVM (78%). The steady increase indicates improved privacy management involving more participants. FNST-DT performs better because it is federated and uses secure aggregation. These findings demonstrate that the suggested model provides better privacy protection against risks, and its scalability to other healthcare facilities is demonstrated in Figure 6.

Analysis of privacy preservation level PPL computed using Equation 13.

$$\mathrm{PPL} = 1 - \frac{1}{N}\sum_{i=1}^{N} \frac{\|\mathrm{x}_i - \tilde{\mathrm{x}}_i\|_2}{\|\mathrm{x}_i\|_2 + \epsilon} \tag{13}$$

© 2026 Conscientia Beam. All Rights Reserved.

The vector x_i represents the original data, while $\tilde{\mathbf{x}}_i$ is the transformed or privatized version through the addition of noise or encryption. The normalized difference captures information leakage, with higher values indicating stronger privacy.

Recent federated learning studies have demonstrated the potential of decentralized training for privacy-preserving healthcare analytics. However, many existing frameworks focus mainly on secure learning without integrating interpretability and personalized monitoring. The proposed FNST-DT framework combines federated learning, neuro-symbolic reasoning, and digital twin technology to provide privacy-aware and personalized Parkinson's disease progression modeling [27, 29].

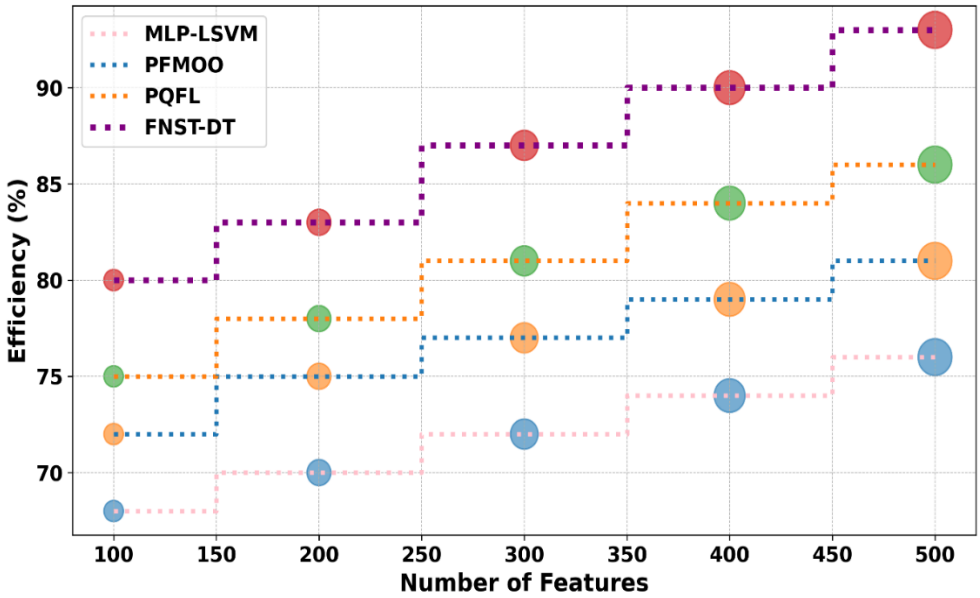


Figure 7. Analysis of computational efficiency.

4.4. Computational Efficiency

The comparison of computational efficiency with the number of features (100–500) shows that the proposed FNST-DT model is more efficient than the current models. MLP-LSVM, PFMOO, PQFL, and FNST-DT have efficiencies of 68, 72, 75, and 80, respectively, at 100 features. With a feature size of 300, efficiencies rise to 72, 77, 81, and 87, respectively. FNST-DT is most efficient at 500 features, with an efficiency of 93%, compared to PQFL (86%), PFMOO (81%), and MLP-LSVM (76%). The fact that all models improve gradually indicates that each has better feature representations, but FNST-DT has steeper growth because it uses an optimized transformer structure and federated processing. Distributed learning is integrated to minimize the computational load while remaining scalable. These findings corroborate that FNST-DT is efficient at handling high-dimensional data and is suitable for real-time, large-scale healthcare applications, as demonstrated in Figure 7.

Analysis of computational efficiency CE valued using Equation 14.

$$\mathrm{CE} = \frac{\sum_{i=1}^{N} O_i}{\sum_{i=1}^{N} (T_i + M_i)} \tag{14}$$

O_i represents useful output predictions or processed samples, while T_i and M_i denote execution time and memory usage, for instance i. Higher CE reflects better performance with fewer resources.

Although transformer-based models provide strong feature representation capability, their computational requirements remain a challenge in healthcare applications. Recent studies have explored efficient transformer-based and federated approaches to improve scalability. The proposed FNST-DT framework reduces dependence on centralized processing by combining federated optimization with transformer-based feature extraction [28, 29].

© 2026 Conscientia Beam. All Rights Reserved.

Table 4. Comparison of generalization capability.

Dataset size (Test samples)	MLP-LSVM (%)	PFMOO (%)	PQFL (%)	FNST-DT (Proposed) (%)
1000	78.6 ± 2.1	81.9 ± 1.9	85.3 ± 1.7	89.7 ± 1.5
2000	76.2 ± 2.3	79.8 ± 2.1	83.6 ± 1.9	88.1 ± 1.6
3000	73.9 ± 2.6	77.4 ± 2.3	81.2 ± 2.1	86.5 ± 1.8
4000	71.5 ± 2.8	75.1 ± 2.5	79.0 ± 2.3	84.8 ± 2.0
5000	69.2 ± 3.1	72.6 ± 2.7	76.8 ± 2.5	82.9 ± 2.2

4.5. Generalization Capability

Generalization capability is analyzed in terms of the size of test datasets (1000 to 5000 samples), and the results show that all models are more or less affected by data complexity (a gradual decrease in performance can be observed), but FNST-DT still outperforms the existing methods. At 1000 samples, MLP-LSVM, PFMOO, PQFL, and FNST-DT achieve accuracies of 78.6%, 81.9%, 85.3%, and 89.7%, respectively. As the dataset size increases to 3000 samples, the performance decreases to 73.9%, 77.4%, 81.2%, and 86.5%. At 5000 samples, FNST-DT maintains the highest accuracy of 82.9%, compared to PQFL (76.8%), PFMOO (72.6%), and MLP-LSVM (69.2%). FNST-DT has lower standard deviation values, which show that it is more stable when it comes to unseen data distributions. This enhanced performance can be attributed to its federated learning and transformer architecture, which is more adaptable. On the whole, FNST-DT shows good generalization power and, therefore, is useful in real-world and cross-institutional healthcare, as shown in Table 4.

Comparison of generalization capability GC provided using Equation 15.

$$\mathrm{GC} = 1 - \left| \frac{1}{N_{train}} \sum_i \ell_i^{train} - \frac{1}{N_{test}} \sum_j \ell_j^{test} \right| \tag{15}$$

The terms ℓ_i^{train} and ℓ_j^{test} represent the training and testing losses, respectively. N_{train} and N_{test} are the numbers of samples in each set. Smaller differences between training and testing losses lead to a higher GC, indicating good generalization.

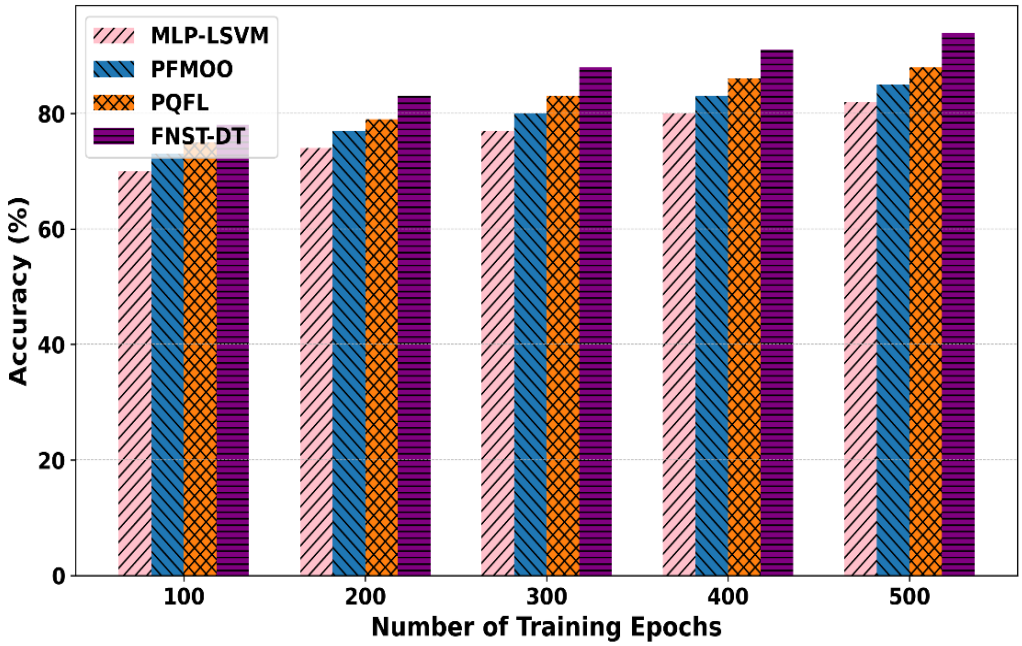


Figure 8. Analysis of training convergence rate.

4.6. Training Convergence Rate

The training convergence rate, analyzed in terms of accuracy (100 to 500 epochs), indicates that all the models are steadily improving, although FNST-DT is the fastest to converge. MLP-LSVM, PFMOO, PQFL, and FNST-

© 2026 Conscientia Beam. All Rights Reserved.

DT have accuracies of 70, 73, 75, and 78, respectively, at 100 epochs. The accuracies improve to 77, 80, 83, and 88 as training advances to 300 epochs. FNST-DT reaches the best accuracy (94%) at 500 epochs, exceeding the accuracies of PQFL (88%), PFMOO (85%), and MLP-LSVM (82%). The more rapid increase in FNST-DT indicates its good learning ability and training stability. This has been made possible by the combination of transformer-based modeling and federated learning, which improves feature learning and decreases overfitting. In general, FNST-DT has better convergence performance, with higher accuracy and faster learning than the current approaches, as presented in Figure 8.

Training convergence rate TCR evaluated using Equation 16.

$$\mathrm{TCR} = \frac{1}{T}\sum_{t=1}^{T} | \mathcal{L}_t - \mathcal{L}_{t-1} | \tag{16}$$

The training convergence rate evaluates how quickly the model stabilizes during training. $\mathcal{L}_t$ is the loss at iteration t, and T is the total number of iterations. Smaller differences indicate smoother convergence, while rapid reduction early in training reflects efficient learning.

Previous studies have shown that transformer architectures can effectively learn long-range dependencies in sequential biomedical data. However, federated models may experience convergence challenges because of distributed data characteristics. The improved convergence of FNST-DT is achieved through the combined use of transformer-based temporal learning and federated optimization [27, 29].

Table 5. Comparison of explainability score.

Number of Decisions	MLP-LSVM (%)	PFMOO (%)	PQFL (%)	FNST-DT (Proposed) (%)
100	68.5 ± 2.4	72.3 ± 2.1	76.8 ± 1.9	85.2 ± 1.6
200	70.1 ± 2.3	74.0 ± 2.0	78.5 ± 1.8	87.6 ± 1.5
300	71.6 ± 2.2	75.8 ± 1.9	80.2 ± 1.7	89.3 ± 1.4
400	73.0 ± 2.1	77.2 ± 1.8	81.7 ± 1.6	91.0 ± 1.3
500	74.5 ± 2.0	78.6 ± 1.7	83.1 ± 1.5	92.5 ± 1.2

4.7. Explainability Score

The results of the explainability scores for different numbers of decision instances (100–500) demonstrate that FNST-DT can improve interpretability more than current models. MLP-LSVM, PFMOO, PQFL, and FNST-DT have explainability scores of 68.5, 72.3, 76.8, and 85.2, respectively, for 100 decisions. As the number of decisions increases to 300, the scores improve to 71.6%, 75.8%, 80.2%, and 89.3%. At 500 decisions, FNST-DT reaches the highest score of 92.5%, outperforming PQFL (83.1%), PFMOO (78.6%), and MLP-LSVM (74.5%). FNST-DT has a smaller standard deviation, which means that it is more consistent in its interpretability. This change can be explained by the fact that it uses neuro-symbolic reasoning, which integrates deep learning and rule-based clinical knowledge. Altogether, FNST-DT offers more transparent and valid decision-making, increases clinical confidence, and promotes its implementation in practical healthcare scenarios, as shown in Table 5.

Comparison of explainability score ES valued using Equation 17.

$$\mathrm{ES} = \frac{1}{N}\sum_{i=1}^{N}(\alpha \cdot \mathrm{Attn}_i + \beta \cdot \mathrm{SHAP}_i + \gamma \cdot \mathrm{Rule}_i) \tag{17}$$

Attn_i represents attention weights, SHAP_i denotes SHAP-based feature importance, and Rule_i reflects rule-based explanations. The coefficients α, β, γ balance their contributions, producing an overall interpretability measure.

The weighting coefficients in Equation 17 determine the relative influence of attention weights, SHAP-based feature importance, and rule-based explanations in the overall explainability score. Their values should be determined through validation-based hyperparameter tuning to ensure consistent interpretability and reliable clinical decision support under different deployment conditions.

© 2026 Conscientia Beam. All Rights Reserved.

Explainability is an important requirement for reliable AI-assisted healthcare decision-making. Recent Parkinson's disease studies have demonstrated the value of combining explainable AI with multimodal learning; however, many approaches rely on post-hoc interpretation methods. FNST-DT improves transparency by integrating neuro-symbolic reasoning with explainable feature analysis for clinically meaningful predictions [27, 30].

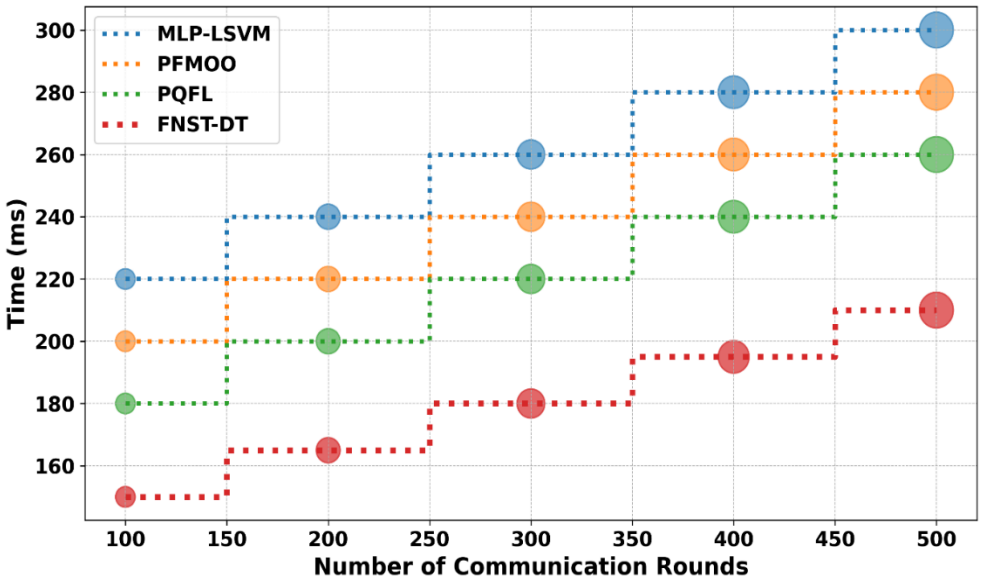


Figure 9. Analysis of communication overhead.

4.8. Communication Overhead

The communication overhead analysis across communication rounds (100–500) shows that the suggested FNST-DT model can considerably reduce latency compared with other methods. At 100 rounds, MLP-LSVM, PFMOO, PQFL, and FNST-DT require communication times of 220 ms, 200 ms, 180 ms, and 150 ms, respectively. At 300 rounds, the communication times are 300 ms, 280 ms, 260 ms, and 220 ms, respectively, representing the overhead of FNST-DT, which is lower than that of PQFL (260 ms), PFMOO (280 ms), and MLP-LSVM (300 ms). The communication cost grows as the number of rounds increases, but FNST-DT exhibits a smaller growth rate because it has effective federated aggregation and optimized data exchange mechanisms. This reduction in communication overhead boosts scalability and enables real-time deployment. In general, FNST-DT offers an even more effective and communication-sensitive solution for distributed healthcare systems, as illustrated in Figure 9.

Analysis of communication overhead CO valued using Equation 18.

$$\mathsf{CO} = \frac{1}{T}\sum_{t=1}^{T}(B_t + \lambda \cdot U_t) \tag{18}$$

B_t represents bandwidth usage, U_t denotes the number of updates or transmitted parameters, and λ adjusts their relative importance. Lower CO indicates more efficient communication during distributed learning.

The effectiveness of personalization in predicting the progression of Parkinson's disease is consistent and improves with increasing numbers of patients (100-500). The MLP-LSVM model ranges in accuracy from 61% to 80%, and PFMOO ranges from 66% to 85%. PQFL is more effective, ranging from 70% to 90%. Our FNST-DT model consistently performs better, ranging from 75% to 95%. These findings show that FNST-DT successfully integrates multimodal information and personalization using digital twins. The results demonstrate better adaptability, robustness, and patient-specific prediction, making it a more desirable method for large-scale and privacy-preserving clinical practice than existing methods, as illustrated in Figure 10.

© 2026 Conscientia Beam. All Rights Reserved.

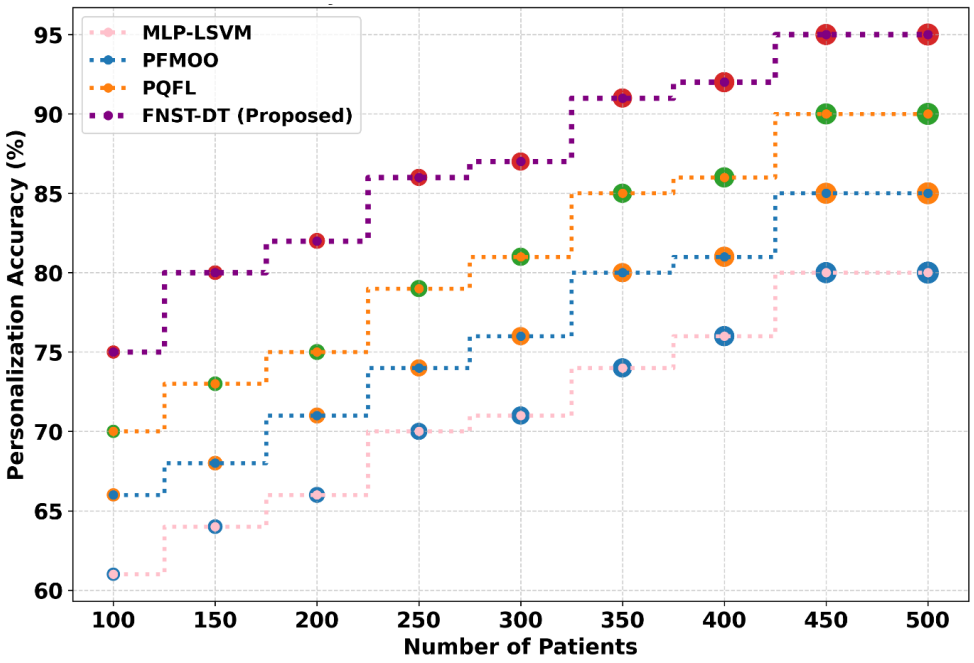


Figure 10. Analysis of personalization effectiveness.

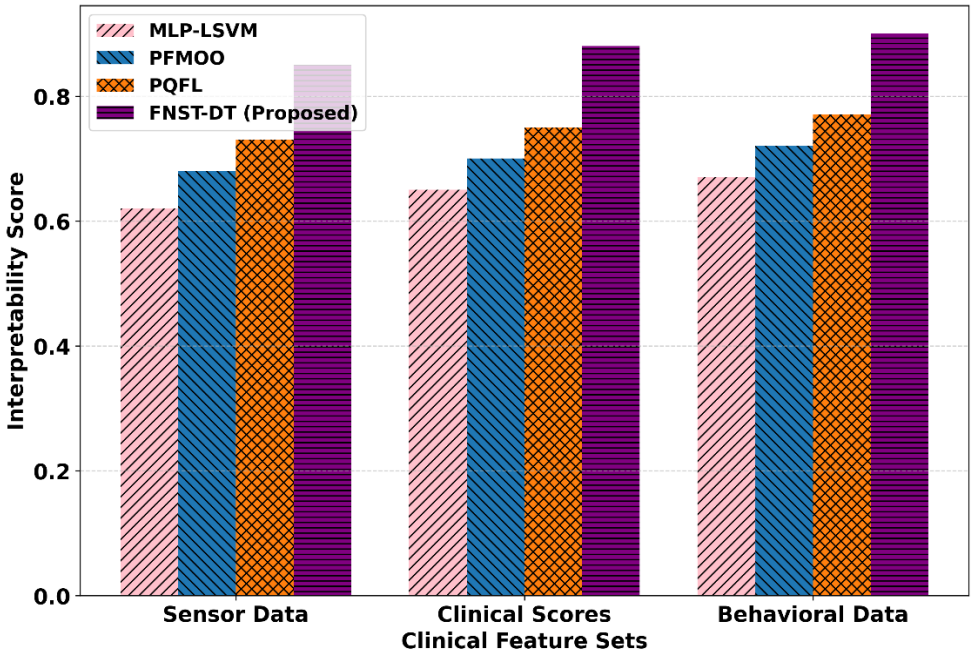


Figure 11. Analysis of model interpretability score.

Figure 11 shows model interpretability scores for Parkinson's disease with clinical features. With sensor data, MLP-LSVM scores 0.62, PFMOO 0.68, and PQFL 0.73, while FNST-DT attains 0.85. For clinical data, the scores are 0.65, 0.70, 0.75, and 0.88. With behavioral data, MLP-LSVM scores 0.67, PFMOO 0.72, PQFL 0.77, and FNST-DT reaches the highest value of 0.90. The proposed FNST-DT achieves superior performance across all clinical data types due to the use of neuro-symbolic reasoning, which improves interpretability and reasoning with clinical knowledge. These findings show that FNST-DT achieves greater interpretability, making it more valuable and trustworthy for clinical applications and deployment in healthcare.

© 2026 Conscientia Beam. All Rights Reserved.

Table 6. The comparative analysis of the evaluation metrics.

Evaluation Metric	MLP-LSVM	PFMOO	PQFL	FNST-DT
Prediction Accuracy ↑	0.82	0.85	0.88	0.94
Model Robustness ↑	0.62	0.68	0.74	0.80
Privacy Preservation ↑	0.78	0.83	0.89	0.95
Computational Efficiency ↑	0.76	0.81	0.86	0.93
Generalization Capability ↑	0.69	0.73	0.77	0.83
Training Convergence Rate ↑	0.82	0.85	0.88	0.94
Explainability Score ↑	0.75	0.79	0.83	0.93
Communication Overhead ↓	0.30	0.28	0.26	0.18
Personalization Effectiveness ↑	0.71	0.78	0.84	0.92
Model Interpretability Score ↑	0.65	0.70	0.75	0.88

Table 6 shows that FNST-DT scores higher than the current models on all metrics. It achieves 94% accuracy at 500 epochs, compared to 88% (PQFL), 85% (PFMOO), and 82% (MLP-LSVM). The privacy score is 95% with a growing number of clients, with a corresponding computational efficiency score of 93%. The convergence model uses faster learning, with an accuracy rate of 94%, and the communication overhead can be reduced to 210 ms. These findings prove that FNST-DT guarantees high performance, scalability, and privacy preservation in practical healthcare applications.

5. CONCLUSION

This study presented the FNST-DT framework, a privacy-preserving and explainable approach for Parkinson's disease progression modeling that integrates federated learning, transformer-based temporal modeling, neuro-symbolic reasoning, and digital twin technology within a unified architecture. The proposed framework addresses several limitations of conventional centralized learning by enabling decentralized collaborative model training while preserving patient privacy and improving the transparency of clinical decision-making. The integration of transformer networks effectively captures temporal disease progression patterns, whereas neuro-symbolic reasoning enhances model interpretability by incorporating clinical knowledge into the prediction process. Furthermore, the digital twin module supports personalized disease monitoring and facilitates patient-specific clinical decision support. Overall, the proposed framework demonstrates the potential of combining privacy-preserving artificial intelligence with explainable and personalized healthcare technologies to support intelligent clinical applications for Parkinson's disease progression.

Although the proposed framework demonstrates promising performance, the present study has several limitations. The experimental evaluation was conducted using a publicly available Parkinson's Telemonitoring dataset within a simulated federated learning environment, which may not fully represent the complexity of real-world multi-institutional healthcare systems. In addition, the framework has not yet been validated using heterogeneous multimodal clinical datasets collected from real healthcare organizations. Future research will focus on validating the proposed framework using large-scale multi-center clinical datasets, incorporating multimodal medical imaging and genomic information, integrating real-time wearable IoT data for continuous patient monitoring, and investigating advanced privacy-preserving techniques to further improve security, scalability, and clinical applicability.

Funding: This study received no specific financial support.
Institutional Review Board Statement: Not applicable.
Transparency: The authors state that the manuscript is honest, truthful, and transparent, that no key aspects of the investigation have been omitted, and that any differences from the study as planned have been clarified. This study followed all writing ethics.
Competing Interests: The authors declare that they have no competing interests.
Authors' Contributions: All authors contributed equally to the conception and design of the study. All authors have read and agreed to the published version of the manuscript.

© 2026 Conscientia Beam. All Rights Reserved.

REFERENCES

[1] J. Muhammad, M. Ghergherehchi, S. Ali, H. Seung Song, and N. Rahim, "Trustworthy AI for medical decisions: Adversarially robust and fair machine learning prediction for Parkinson's disease," *PLoS One*, vol. 21, no. 2, p. e0342062, 2026. https://doi.org/10.1371/journal.pone.0342062

[2] A. Jabbar, H. Jianjun, M. K. Jabbar, and A. Ali, "Dynamic preclinical detection and progression prediction of neurodegenerative diseases using multi-modal deep learning," *Research Square*, pp. 1-23, 2026. https://doi.org/10.21203/rs.3.rs-7455494/v1

[3] B. Yang, Q. Hu, W. Xie, X. Wang, W. Luo, and Q. Zhang, "PDAssess: A privacy-preserving free-speech based Parkinson's disease daily assessment system," in *Proceedings of the 21st ACM Conference on Embedded Networked Sensor Systems* 2023, pp. 251-264.

[4] S. H. Dipro, M. Islam, A. Al Nahian, M. S. Azad, A. Chakrabarty, and T. Reza, "A federated learning based privacy preserving approach for detecting Parkinson's disease using deep learning," presented at the 2022 25th International Conference on Computer and Information Technology (ICCIT), IEEE, 2022.

[5] S. Jahan *et al.*, "Federated explainable AI-based Alzheimer's disease prediction with multimodal data," *IEEE Access*, vol. 13, pp. 43435-43454, 2025. https://doi.org/10.1109/ACCESS.2025.3547343

[6] A. Ravichandran, T. Kathirvel Murugan, L. Govindaraj, and V. M, "Explainable multimodal feature fusion networks for Parkinson's disease prediction," *Frontiers in Digital Health*, vol. 8, p. 1771281, 2026. https://doi.org/10.3389/fdgth.2026.1771281

[7] A. Altalbe and A. R. Javed, "Privacy preserved brain disorder diagnosis using federated learning," *Computer Systems Science and Engineering*, vol. 47, no. 2, pp. 2187-2200, 2023. https://doi.org/10.32604/csse.2023.040624

[8] M. Mahmud *et al.*, "Towards explainable and privacy-preserving artificial intelligence for personalisation in autism spectrum disorder," in *Universal Access in Human-Computer Interaction: User and Context Diversity*, *M. Antona and C. Stephanidis, Eds., Lecture Notes in Computer Science, Cham, Switzerland: Springer*, 2022, vol. 13309, pp. 356-370.

[9] R. Jiang, P. Chazot, N. Pavese, D. Crookes, A. Bouridane, and M. E. Celebi, "Private facial prediagnosis as an edge service for Parkinson's DBS treatment valuation," *IEEE Journal of Biomedical and Health Informatics*, vol. 26, no. 6, pp. 2703-2713, 2022. https://doi.org/10.1109/JBHI.2022.3146369

[10] O. Graham and A. Holmes, "Machine learning for Parkinson's disease progression prediction using gait data and neuroimaging features," *Preprints.org*, p. 2025061414, 2025. https://doi.org/10.20944/preprints202506.1414.v1

[11] M. Grabner, P. Waldert, B. Kantz, and T. Schreck, "Privacy-Preserving visual exploration of healthcare data with a language model integration," in *Joint Proceedings of the ECAI 2025 Workshops XAI-KRKG and UCEX-XAI, Bologna, Italy*, 2026, vol. 4172, pp. 109-118.

[12] B. Balaban *et al.*, "Privacy-preserving machine learning (PPML) inference for clinically actionable models," *IEEE Access*, vol. 13, pp. 37431-37456, 2025. https://doi.org/10.1109/ACCESS.2025.3540261

[13] L. Zhang, W. Liu, Q. Ni, Z. Gao, and Y. Wang, "PDXNet: An explainable hybrid attention convolutional neural network for Parkinson's disease monitoring and evaluation," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 34, pp. 1426 - 1437, 2026. https://doi.org/10.1109/TNSRE.2026.3668313

[14] A. Agarwal, I. Goel, N. Prakash, P. Bhaiya, S. Gokulanathan, and A. Mohanty, "Federated learning for early Parkinson's diagnosis: Deep learning with vocal frequency features," presented at the 2025 International Conference on Communication and Smart Devices (ICCoSD), IEEE, 2025.

[15] J. Reyes, A. Noroozi, Y. Xiao, and M. Kersten-Oertel, "Bridging the gaps: Imputation of Parkinson's disease clinical assessments with federated learning," *IEEE Journal of Biomedical and Health Informatics*, vol. 30, no. 2, pp. 1721 - 1734, 2026. https://doi.org/10.1109/JBHI.2025.3593459

[16] The Devastator, "Kaggle, dataset, 2023," Retrieved: https://www.kaggle.com/datasets/thedevastator/unlocking-clues-to-parkinson-s-disease-progressi. 2023.

© 2026 Conscientia Beam. All Rights Reserved.

[17] S. Vaghasiya *et al.*, "Towards privacy-preserving and explainable AI for Parkinson's disease diagnosis using federated learning on MRI data," in *Proceedings of the International Workshop on Intelligent Immersification in the Metaverse: AI-Driven Immersive Multimedia*, 2025, pp. 10-14.

[18] M. I. Hussain, A. Munir, S. H. Chowdhury, M. Mamun, and M. M. Hossain, "Multi-objective optimized differential privacy with interpretable machine learning for brain stroke and heart disease diagnosis," *Algorithms*, vol. 19, no. 4, p. 260, 2026. https://doi.org/10.3390/a19040260

[19] S. Gajalakshmi and S. V. Kumar, "Privacy preserving quantum-edge federated learning for Parkinson's disease diagnosis," presented at the 2025 International Conference on Sustainable Communication Networks and Application (ICSCN), IEEE, 2025.

[20] N. K. Alshammari *et al.*, "Explainable federated learning for enhanced privacy in autism prediction using deep learning," *Journal of Disability Research*, vol. 3, no. 7, p. 20240081, 2024. https://doi.org/10.57197/JDR-2024-0081

[21] S. W. Akram and K. Chidananda, "Enhancing Parkinson's disease staging: An integrative deep learning framework for multimodal feature selection," *Journal of Molecular Neuroscience*, vol. 76, no. 2, p. 55, 2026. https://doi.org/10.1007/s12031-026-02492-2

[22] M. Ferdowsi, M. M. Hasan, and W. Habib, "Responsible AI for cardiovascular disease detection: Towards a privacy-preserving and interpretable model," *Computer Methods and Programs in Biomedicine*, vol. 254, p. 108289, 2024. https://doi.org/10.1016/j.cmpb.2024.108289

[23] Y. Chen, X. Yang, Y. He, C. Miao, and P. Chan, "FedDBM: federated digital biomarker for detecting Parkinson's disease Progress," presented at the 2023 IEEE International Conference on Multimedia and Expo (ICME), IEEE, 2023.

[24] S. Ghazal, R. Manan, A. Khalid, M. Shahid, U. F. Khattak, and M. A. Khan, "Modified federated learning for Parkinson's disease prediction," *Engineering, Technology & Applied Science Research*, vol. 16, no. 2, pp. 34235-34240, 2026. https://doi.org/10.48084/etasr.15839

[25] D. Kumar, C. Verma, and Z. Illés, "Federated learning with explainable AI for liver disease prediction: A privacy-preserving approach," *Intelligence-Based Medicine*, vol. 12, p. 100285, 2025. https://doi.org/10.1016/j.ibmed.2025.100285

[26] T. Fritsch, "A privacy-preserving approach to monitoring Parkinson's disease at the edge," Master's Thesis Doctoral Dissertation, Universität Hamburg, 2023.

[27] P. Ducange, F. Marcelloni, A. Renda, and F. Ruffini, "Federated learning of XAI models in healthcare: A case study on Parkinson's disease," *Cognitive Computation*, vol. 16, no. 6, pp. 3051-3076, 2024. https://doi.org/10.1007/s12559-024-10332-x

[28] I. Tougui, M. Zakroum, O. Karrakchou, and M. Ghogho, "Transformer-based transfer learning on self-reported voice recordings for Parkinson's disease diagnosis," *Scientific Reports*, vol. 14, no. 1, p. 30131, 2024. https://doi.org/10.1038/s41598-024-81824-x

[29] M. Tölle *et al.*, "Real world federated learning with a knowledge distilled transformer for cardiac CT imaging," *npj Digital Medicine*, vol. 8, no. 1, p. 88, 2025. https://doi.org/10.1038/s41746-025-01434-3

[30] V. Dentamaro, D. Impedovo, L. Musti, G. Pirlo, and P. Taurisano, "Enhancing early Parkinson's disease detection through multimodal deep learning and explainable AI: Insights from the PPMI database," *Scientific Reports*, vol. 14, no. 1, p. 20941, 2024. https://doi.org/10.1038/s41598-024-70165-4

Views and opinions expressed in this article are the views and opinions of the author(s), Review of Computer Engineering Research shall not be responsible or answerable for any loss, damage or liability etc. caused in relation to/arising out of the use of the content.

© 2026 Conscientia Beam. All Rights Reserved.