Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/IN05.20250198
Published: November 6, 2025

A Sequential and Progressive Multimodal Machine Learning Framework for Risk Assessment of Parkinson's Disease Using PPMI Data

Manjula R
Professor, School of Computer Science and Engineering
Vellore Institute of Technology, Vellore
Tamil Nadu, India
rmanjula@vit.ac.in

Nandini H. Sarvaiya, Harshitaa Kashyap and Helan Maria Cyrill
School of Computer Science and Engineering
Vellore Institute of Technology, Vellore
Tamil Nadu, India
sarvaiyanandini.h2022@vitstudent.ac.in, harshitaa.kashyap2022@vitstudent.ac.in,
helanmaria.cyrill2022@vitstudent.ac.in

Abstract

Parkinson's disease is a progressive brain disorder that lowers dopamine levels and leads to both movement and non-movement difficulties. Because its early signs are subtle and the course of the disease can vary widely between patients, diagnosis and monitoring remain challenging. In this study, we present a machine learning framework designed to improve risk assessment by using multimodal data from the Parkinson's Progression Markers Initiative (PPMI). At the core of our approach is a sequential and progressive model that integrates motor and non-motor clinical features with information about how symptoms change over time, enabling more reliable prediction of both disease risk and trajectory across diverse patient groups. The framework's performance is evaluated with standard measures of diagnostic accuracy, sensitivity, and progression forecasting. A review of existing work shows that fewer than 15% of machine learning studies on Parkinson's use multimodal data in depth, and fewer than 5% combine sequential symptom tracking with progression risk scoring. This highlights the novelty of our approach. Our findings shed light on which features are most important, demonstrate the potential of longitudinal tracking for emerging symptoms, and suggest how such models can support earlier detection and more personalized care. We also discuss practical factors for applying clinical data in scalable and reliable healthcare AI systems. Overall, this work provides insights that can guide future research and inform best practices in using machine learning for Parkinson's disease management.

Keywords
Parkinson's disease, Machine learning, Multimodal data, Disease progression, Early detection

1. Introduction
Parkinson's disease (PD) is one of the fastest-growing neurodegenerative disorders in the world today, affecting over ten million people. It is characterized by both motor and non-motor symptoms that result from gradual loss of neurons in several regions of the brain, especially those involved in movement and regulation of neurotransmitters like dopamine, serotonin, and norepinephrine. The way that these symptoms appear is quite complex and heterogeneous in nature, making early diagnosis and accurate monitoring very difficult. Current clinical assessments are usually made

periodically such as the Unified Parkinson's Disease Rating Scale (UPDRS), which can provide only a high-level observation and usually fails to capture the more subtle symptom fluctuations that occur over time.

Recent advances in digital health and artificial intelligence have created new opportunities for objective, data-driven assessment of PD. Large-scale initiatives such as the Parkinson's Progression Markers Initiative (PPMI) have made multimodal datasets publicly available, containing motor, cognitive, imaging, and biospecimen information. However, despite these resources, most existing predictive models still focus on a single-modality input, which fails to give an output that can be broadly applied or that can reliably capture how symptoms evolve. Even when they do, they fail to integrate both motor and cognitive symptom trajectories into a unified risk-assessment framework that is temporally aware. This gap in research hinders our understanding of how PD evolves and thus negatively affects the chances of earlier intervention through development of personalized monitoring tools.

The motivation for this research lies within the need to design clinically reliable, temporally aware machine-learning systems that can capture complex, multimodal patterns in PD progression. The core problem addressed in this work is the lack of integrative frameworks that can combine heterogeneous patient data sources and represent the temporal progression of symptoms in a scalable, interpretable way. In order to close these gaps, this paper proposes a multimodal learning pipeline using clinical data from the PPMI cohort. It does so by analyzing the interactions between different modalities over time as well as their individual progression over time. The ultimate aim is to enhance accuracy of diagnosis, continuously track a variety of symptoms, and create a more personalized understanding of Parkinson's progression.

1.1 Objectives

This research takes important steps towards addressing the main drawbacks of existing studies on Parkinson's disease through focusing on multimodality and temporal progression. It proposes to address the limitations of single-modality studies and enhance the understanding of real-world clinical applicability of data by combining several heterogeneous data sources, including motor, imaging and non-motor clinical data, to model the progression of Parkinson's disease (PD) more comprehensively. The research also focuses on characterizing temporal dynamics of symptom progression, by developing models that represent symptom progression over time using longitudinal data. The focus on temporal dynamics feeds into another goal of the research, which is to build multimodal learning pipelines and advanced sequence encoders (e.g., RNNs, LSTMs, and CNNs) to learn cross-domain interactions and multi-dimensional and complex temporal patterns in longitudinal data. The research will demonstrate model performance and model interpretability on several encoder-fusion architectures, while identifying clinically relevant and interpretable models of progression. It will begin the development of "personalized patient profiles" through the production of interpretable and individualized progression trajectories for early detection monitoring, as well as targeted and customized individualized intervention planning. Ultimately, the research aims to provide robust, scalable, and interpretable frameworks for modeling Parkinson's disease (PD) progression to assist with single-modality studies that result in very limited applicability related to real-world clinical practice.

2. Literature Review

Parkinson's disease (PD) is a progressive neurodegenerative disorder impacting millions of individuals globally, characterized by both motor symptoms like tremor, rigidity and bradykinesia and non-motor symptoms such as cognitive decline, mood disorders and sleep disturbances. Traditional clinical assessments, such as the Unified Parkinson's Disease Rating Scale (UPDRS), can provide substantial insights but they fail to capture the dynamic and heterogeneous nature of PD symptom progression over time. Hence, early and accurate detection continues to be a key limitation in PD management rendering it difficult to intervene during the early stage which could slow down the pace of the progression of disease significantly and thus, help improve quality of patient care.

2.1 Motor Symptom Assessment Using Wearable Sensors

Due to recent advances in the field of wearable sensor technology, the monitoring of motor symptoms in PD has improved significantly. Sotirakis et al. (2023) performed an analysis with 74 PD patients over 18 months using six Inertial Measurement Units (IMUs) which recorded walking and postural sway. A Random Forest-based machine learning model was used which estimated MDS-UPDRS-III scores with a root mean square error (RMSE) of 10.02 and detected statistically significant motor symptom progression over the course of 15 months. In this same timeframe, traditional clinical rating scales failed to capture any measurable change. Zhang et al. (2024) tested wearable sensor-based gait analysis devices in PD, demonstrating that gait impairments worsen over time and can be numerically

quantified using accelerometers and gyroscopes. A review by Rodríguez-Martín et al. (2024) evaluated commercial symptom monitoring devices including PKG™, Kinesia 360™, and PDMonitor™.This helped discover their potential for continuous motor symptom tracking outside clinical settings as well as found critical limitations such as variable correlation with clinical observations, usability challenges affecting patient adherence, and lack of independent validation studies. These findings show that wearable technology can provide effective results, but the standardization and the robust validation are essential for clinical adoption (Sotirakis et al. 2023, Zhang et al. 2024, Rodríguez-Martín et al. 2024).

2.2 Non-Motor Symptom Assessment and Clinical Impact

Non-motor symptoms such as cognitive decline, mood disorders, autonomic dysfunction, and sleep disturbances have been recognized as major contributors to PD and often precede the appearance of motor symptoms by years. A Non-Motor Symptoms Scale (NMSS) was developed and validated by Chaudhuri et al. (2007) which used 30-item rater-based instrument assessment of nine domains of non-motor symptoms. It has been utilized in over 100 clinical tests and has repeatedly shown strong correlations between non-motor symptom level and measures of health-related quality of life (van Wamelen et al. 2021). A cross-sectional analysis was done that examined 1,607 PD patients across Europe, the Americas, and Asia which revealed a mean NMSS score of 46.7 ± 37.2, with significant geographical variations in level of non-motor symptoms appearance. Mood/apathy domain emerged as the strongest predictor of the quality of life ($\beta = 0.308$, $p < 0.001$) (van Wamelen et al. 2021). Even though non-motor symptoms have a substantial effect on the patients' quality of life and they have potential as early diagnostic tools, they still remain underrepresented in routine clinical evaluations, leading to delayed diagnosis and inadequate management strategies (Chaudhuri et al. 2007, van Wamelen et al. 2021).

2.3 Multimodal Data Integration and Advanced Fusion Techniques

Multiple modalities together such as motor assessments, cognitive evaluations, neuroimaging, and molecular biomarkers provide a more complete and holistic view of PD's heterogeneous nature. Zhou et al. (2025) developed a novel multimodal classifier using data from the Parkinson's Progression Markers Initiative (PPMI). This includes hematological information, proteomics, RNA sequencing, metabolomics, and dopamine transporter (DaTscan) imaging. This helped achieve a balanced classification accuracy of 97.7% through a multimodal encoder-based model that also had multi-head cross-attention (MMT_CA). The novelty in this above-mentioned work lies in the cross-attention mechanism's ability to identify the complex inter-dependencies across many different data modalities. This is what helped it to significantly outperform the traditional approaches. SHAP (SHapley Additive exPlanations) analysis recognized the most important biomarkers for diagnosis across modalities to provide interpretability which is necessary for clinical trust and precision medicine applications. Benredjem et al. (2024) introduced the Parkinson Multi-Modal Prediction Framework (PMMD), which integrated handwriting images, spiral drawings, and clinical data. A cross-modal attention architecture was used which helped achieve 96% accuracy on independent test sets and surpassed previous state-of-the-art approaches. Even after benefitting from these advances, challenges persist in effectively merging heterogeneous data types. Managing the dimensional disparities between high-dimensional imaging data and lower-dimensional clinical metadata is the most difficult of them all. (Zhou et al. 2025, Benredjem et al. 2024).

2.4 Longitudinal Analysis and Temporal Progression Modelling

Even though longitudinal research is very important for capturing PD's evolving nature, temporal modeling remains underdeveloped in current literature. A review examining machine learning models for predicting PD progression using longitudinal data found that Long Short-Term Memory (LSTM) networks and different ensemble methods like Random Forest or Light Gradient Boosting Machine (LGBM) are the most effective at measuring disease progression accurately. LSTM models achieved an accuracy of up to 90% and AUC scores of 93.79%, while LGBM models showed 90.73% accuracy and AUC of 94.57%. Hu et al. (2025) addressed the lack of temporal dependency modeling by proposing a unified framework which integrates graph neural networks (GNNs) to model structural relationships between different clinical symptoms and Transformer architectures. This helped recognize the dynamic temporal features during disease progression. A structure-aware gating mechanism aided by dynamically adjusting fusion weights between structural encodings and temporal features. Mellema et al. (2024) used functional MRI (fMRI) connectivity patterns to predict progression in MDS-UPDRS-III and Montreal Cognitive Assessment (MoCA) scores over one year. This led to the achievement of a state-of-the-art performance with MAE of 1.8 and R^2 of 0.69 for motor progression as well as recognized the connectivity between deep nuclei, motor regions, and the thalamus as most

predictive. Even with these advances, most models still treat disease progression as a static classification problem rather than sequential evolution of symptoms over time. (Hu et al. 2025, Mellema et al. 2024).

2.5 Research Gaps and Future Directions

Even though a lot of progress has been made, several gaps remain in current PD research. Most multimodal approaches perform one-time static classification rather than modelling the sequential appearance of symptoms as the disease progresses. This leads to missed opportunities for stage-specific interventions. Existing models are typically trained on fixed datasets and do not incorporate the dynamic updating of quantified risk predictions as new symptom data becomes available over a patient's treatment timeline. Many high-performing models are trained or validated on homogeneous, single-centre datasets which limits their applicability across different populations, geographical regions, and clinical settings. Despite the accuracy being reported high, practical challenges remain including handling missing modalities, managing variability in data quality and ensuring easy interpretation for use in a clinical decision-making setting. Finally, while non-motor symptoms often precede appearance of motor symptoms and strongly correlate with quality of life, they remain underrepresented in predictive models and clinical assessments (Zhou et al. 2025, Hu et al. 2025, van Wamelen et al. 2021).

2.6 Summary

Literature shows that multimodal machine learning approaches, i.e. approaches that include diverse data sources including wearable sensors, neuroimaging, clinical assessments, and molecular biomarkers, significantly outperform single-modality models in PD detection and classification. Advanced fusion techniques, particularly attention-based architecture, can accurately recognize and map cross-modal dependencies which improve diagnostic accuracy. Longitudinal and temporal modeling can track disease progression with greater sensitivity than traditional clinical scales. Ultimately, critical gaps do remain in sequential symptom progression modeling, real-time updation of predicted risk scores, diversity of datasets, and feasibility of deploying such a model in a clinical setting. The following proposed research addresses these limitations by developing a machine learning framework that integrates multimodal data, models temporal symptom changes, supports dynamic risk updating, and provides interpretable outputs for PD risk assessment and longitudinal symptom tracking.

3. Methods

The suggested framework presents a multimodal deep learning pipeline that robustly and transparently predicts disease stage in Parkinson's Disease (PD). The end-to-end system employs three complimentary data modalities to describe the heterogeneous continuum of PD– motor symptoms, non-motor symptoms, and imaging features. The modalities undergo modality-specific preprocessing and coding prior to their combination using an intermediate fusion mechanism, ultimately providing a predicted disease stage (early, mid or late) and confidence scores. The pipeline employs variational autoencoders (VAEs) such that likely compact and noise resilient latent representations are learned, with the goals of improving generalizability and robustness to data which may be incomplete or contain noise.

3.1 Pipeline Overview

The global pipeline begins with the intake of multimodal raw data, which includes both motor assessments (for instance, Unified Parkinson's Disease Rating Scale [UPDRS] scores, gait-related, etc.) and non-motor assessments (again, cognitive, psychological, sleep, autonomic, sensory, etc.) possibly including MRI and DTI imaging data. Each data stream is pre-processed in isolation to handle issues such as missing data, normalization, and alignment by time. After preprocessing, modality-specific encoders output latent representations which characterize the latent structure and variability of each data type, these representations are then fused in a common latent space followed by a classifier to output predicted stage of disease.

Thus, the architecture progresses through a series of steps as outlined above: Raw data → Modality-Specific encoders → Fusion Layer → Classifier → Prediction Output. The VAEs (variational autoencoders), used as part of the architecture, allow for a better potential for learned latent representations to be smooth, continuous, and overfitting resistant, due to the introduction of a probabilistic regularizer. This provides the model with the opportunity to make sturdy predictions, even when some, or all, of the relevant data has been lost or are only partially remaining, as is often the case in clinical datasets, (for instance, PPMI).

3.2 Modality-Specific Encoder Pipelines

Each data modality is processed through an encoder specifically tailored to its statistical and structural properties.

Processing of Motor Data: Motor-related data, which includes time-series evaluations of movement or gait, is organized into a temporally aligned structure and pre-processed by means of forward and backward imputation to fill the gaps. The modality uses a transformer-based encoder for their ability to model long-range temporal dependencies and inter-relationships among features based on attention. The output is an efficient 16-dimensional latent vector that captures the patient's motor performance over time.

Processing of Non-Motor Data: Non-motor features consist of several symptom domains of cognitive, psychological, sleep, autonomic and sensory measures and thus, each measure captures a different domain of PD pathology. We first encoded the data categorically (where applicable), normalized them to be on similar scales, and performed feature selection for each non-motor symptom domain. We implemented non-motor symptom domain-specific multilayer perceptron (MLP)-based variational encoders, for each differentiator of non-motor symptom domain. The cognitive (101 features), psychological (35), sleep (13), autonomic (35), and sensory (21) encoders all project their features into a unified 16-dimensional latent space, that provides a balanced spread of the different types of symptoms while still preserving their unique characteristics.

Imaging Data Processing: Imaging characteristics are obtained from structural MRI and DTI scans designed to measure morphologic and microstructural characteristics of the brain. MRI characteristics consist of a composite of 115 site-specific metrics sampled from both the cortical and subcortical aspects of the brain, while DTI measures consisted of a total of 32 metrics focusing on white matter tracts. Before encoding, each of the imaging features were standardized (z-score normalization), missing values were imputed using a median strategy. Each of the imaging modalities were then put through the variational encoder which is based on Convolutional Neural Networks (CNN) and learns a spatially structured representation of anatomical variation. The output from the encoder supplies a 16-dimensional latent embedding that represents high-level features of the imaging data in an interpretable and parsimonious form.

3.3 Fusion/Integration Layer

Because PD data is multimodal in nature, an effective way of combining the feature spaces is necessary. Here we consider an intermediate fusion strategy, specifically in our intermediate fusion strategy we concatenate the latent representations from each modality-specific encoder. Each latent representation is a 16-dimensional vector: one from motor, one from non-motor, and one from imaging. The result is a fused vector representation of 48 dimensions, which is passed onto a shared multilayer perceptron (MLP) classifier that maps the feature space to the final predicted disease stage.

Merging modalities at an intermediate level strikes a balance between interpretability and richness in interactions (e.g. late fusion vs early fusion). The intermediate fusion design permits the model to combine across modalities (e.g. it can capture some relationship between motor impairments and change in neuroanatomy), while not demanding a complete understanding (an interpretation of all modalities together, from an input level). The design also allows each encoder to specialize in its data type before we combine the modalities.

In terms of empirical evaluation, the intermediate fusion approach reached a cross-validation accuracy of 89.5% in the PPMI dataset--a performance level on par with late fusion while utilizing fewer parameters and posing a reduced risk of overfitting due to PPMI's modest sample size. The intermediate fusion also has the advantage of resilience to missing or noisy modalities, for example, 20-30% missing DTI data, since each encoder can each model their own input space independently. Recent comparative work in multimodal PD analysis cites the intermediate fusion approach with similar architectures achieving almost 90% accuracy (PMMD 2024). In general, the intermediate fusion model is a practical integration approach that is conceptually transparent and a good choice for researchers that want assurance of reliability while taking implementation simplicity into account.

3.4 Model Architecture and Algorithms

The fundamental component of architecture flows as follows: Input → Modality Encoders → Intermediate Fusion → Classifier → Output. Each encoder is represented as a variational autoencoder (VAE), containing a fully connected encoder network, a probabilistic latent layer, and a decoder (used during training for reconstruction regularization).

The latent space is arranged as a 64-dimensional representation, from which 16-dimensional modality-specific embeddings can be extracted for fusion.

The VAEs employ the reparameterization trick, defined as $z = \mu + \sigma \odot \epsilon$, where $\epsilon \sim N(0, I)$, to enable backpropagation through the stochastic latent variables. Regularization is applied through Kullback-Leibler (KL) divergence, encouraging the latent space distribution $q(z \mid x)$to approximate a standard normal prior $p(z)$. The KL divergence term is given by: $KL(q(z \mid x) \| p(z)) = -1/2E[1 + \log(\sigma^2) - \mu^2 - \sigma^2]$. The overall training loss combines reconstruction or classification loss $L_{CE}$with the KL term: $L_{Total} = L_{CE} + \beta L_{KL}$ where β=0.1 balances reconstruction fidelity and latent regularization. Regularization layers such as Batch Normalization and Dropout (rate = 0.4) are used to improve stability and prevent overfitting. The classifier consists of two fully connected layers with ReLU activations, mapping the 48-dimensional fused feature vector to the final SoftMax output over three disease stages.

3.5 Training Strategy

Training was done using the Adam optimizer with a learning rate of 10^{-3} and a batch size of 16. A weighted cross-entropy loss was employed to tackle the data imbalance across classes, whereby the model assigned more weight to minority classes. The data was split with an 80/20 stratified split so class proportions were preserved in training and validation folders. In, the course of training, KL divergence regularization (weight = 0.1) kept the latent representations stable, and early stopping was effective to reduce overfitting. The validation accuracy was levelled at approximately 88 to 89.5%, the macro F1-score was 91%, and area under the curve (AUC) reached 0.92. The results also indicated good generalization potential at the class level throughout the three phases (Early: F1= 94.5% Mid: F1 = 85.5%, Late: F1 = 91.5%).

3.6 Hyperparameter Optimization

Hyperparameter tuning was accomplished with Bayesian Optimization using a Gaussian Process surrogate model. This approach allows for the efficiency of balance between exploration (searching unexplored regions of the space), and exploitation (further improvements on already identified good configurations) - accomplished using the Expected Improvement (EI) acquisition function. The search space consisted of learning rate (10^{-4} – 10^{-2}), dropout rate (0.2 - 0.6), batch size, number of hidden units, and latent dimensionality. The evaluation budget was from 50 - 100 trials to be computationally reasonable and have sufficient batches to yield converged model performance.

3.7 Justification of Design Choices

The intermediate fusion approach was chosen because of its clear advantages in balancing performance, interpretability, and ease of implementation. The model can learn from cross-modal information but allows the different representations of the data to remain separate, which is useful given that not all modalities may be available in medical contexts. We incorporated regularization in the latent spaces using variational encoders, which account for the uncertainty present in medical data, and allow for generalization using the probabilistic prior distributions. Lastly, we opted for separate model architectures to remain consistent with the statistical or data structure of each modality (e.g., transformers for temporal motor data, MLPs for tabular non-motor data, CNNs for spatial imaging data). Together, these models create a cohesive and robust framework for multimodal learning of PD progression (Figure 1).

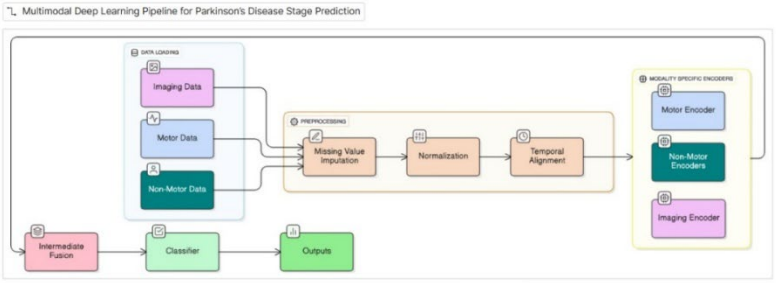

Figure 1. Multimodal learning Pipeline for Parkinson's Disease Prediction

4. Data Collection

The research adopts a complete multimodal method which draws on Parkinson's Progression Markers Initiative (PPMI) database information. The PPMI database contains well-organized patient information that tracks clinical and neuropsychiatric and neuroimaging data in a longitudinal manner. The research depends on three main data types which include motor information and non-motor information and imaging data. The combination of motor and non-motor data with imaging results provides a complete understanding of Parkinson's disease by showing all clinical and biological variations. The multimodal approach reveals intricate relationships between different variables and predictive indicators which single-method studies fail to detect.

The extraction of *non-motor features* occurs through five separate CSV files which contain information about two cognitive subdomains as well as neurological and mood and sleep assessments. These assessments are organized into three time-based categories which include early, mid, and late stages based on the patient's visit identification number. The detailed time-based structure of non-motor symptoms enables researchers to detect subtle differences between patients while tracking how their conditions develop over time. The time-based and detailed structure of non-motor symptoms enables researchers to detect subtle differences between patients while tracking their disease progression. The standardized clinical questionnaires with 100 to 150 patient entries serve as the main source for *motor symptom* information collection. These devices use two main sensors which are axivity and opals. The Opals system uses high-frequency accelerometer data together with gyroscope and magnetometer sensors to achieve accurate measurement of stride length and cadence speed and postural transitions. The axivity device records triaxial movement with high precision which enables continuous motion monitoring and sleep tracking and walking pattern analysis. Our multimodal encoding system uses motor scores as lower priority because these scores have limited dimensions and broad categories which result in weak predictive power for tracking disease progression over time. *Neuroimaging* data collection involves the combination of Magnetic Resonance Imaging (MRI) with Diffusion Tensor Imaging (DTI) techniques. MRI generates high-resolution anatomical images which researchers utilize to evaluate brain region dimensions and cortical thickness for monitoring neurodegenerative progression. DTI allows researchers to study white matter microstructure and brain connectivity through water molecule diffusion measurements which identify early fibre tract changes before patients show clinical symptoms.

Large sample size, standardized data collection methods, and an extended follow-up period that permits in-depth data analysis are all advantages of the PPMI cohort. However, because different study sites employ different data collection techniques, researchers encounter difficulties due to missing data and study participants quitting.

4.1 Data Preprocessing and Cleaning

To ensure consistency, quality, and effective and significant extraction of features from each modality of data, all modalities underwent stringent preprocessing processes. For all the symptom and imaging data, to merge the data, both clinical identifiers and visit identifiers were converted to a comparable string format. Missing values were dealt with by either the imputation of values or excluding columns when there were few missing values. Variables were also brought to a standardized meaning (e.g. z-scores) in order to make the variables more comparable across visits and participants. Time was also accounted for in terms of progressing over time by categorizing visits (and clinical weeks) into early, mid, or late progression phases, and only visits assigned a clinical diagnosis of Parkinson's disease were included in the analyses.

Cognition and Motor Symptoms: The cognitive and motor symptoms were appended, and the columns of symptoms were changed to numerical format. In constructing my variables, I limited the symptoms to a numeric response (no categorical, no binary), and I calculated the mean symptom score per visit, captured symptom peaks, and summarized symptom changes in table format to show general change progress.
Non-Motor Symptoms: Several non-motor symptoms CSVs, including psychiatric, sleep, and autonomic domains, were utilized. Binary indicators were created for the presence or absence of symptomology, with presence averages computed for each visit.

Motor Symptoms: Motor function and quality-of-life CSVs were combined to calculate total motor scores, and motor disease stages were assigned based upon existing clinical cutoff thresholds. A combined motor stage label was developed to facilitate the longitudinal tracking.

Imaging: The MRI and DTI attributes CSV files were merged, and the attributes chosen addressed anatomical measurements (e.g., cortical thickness, regional brain volumes) and image quality metrics (e.g., entropy, contrast-to-noise ratio). Scans underwent quality control, and then the disease stage was inferred using either clustering algorithms or threshold-based mapping of the imaging markers (Figure 2).

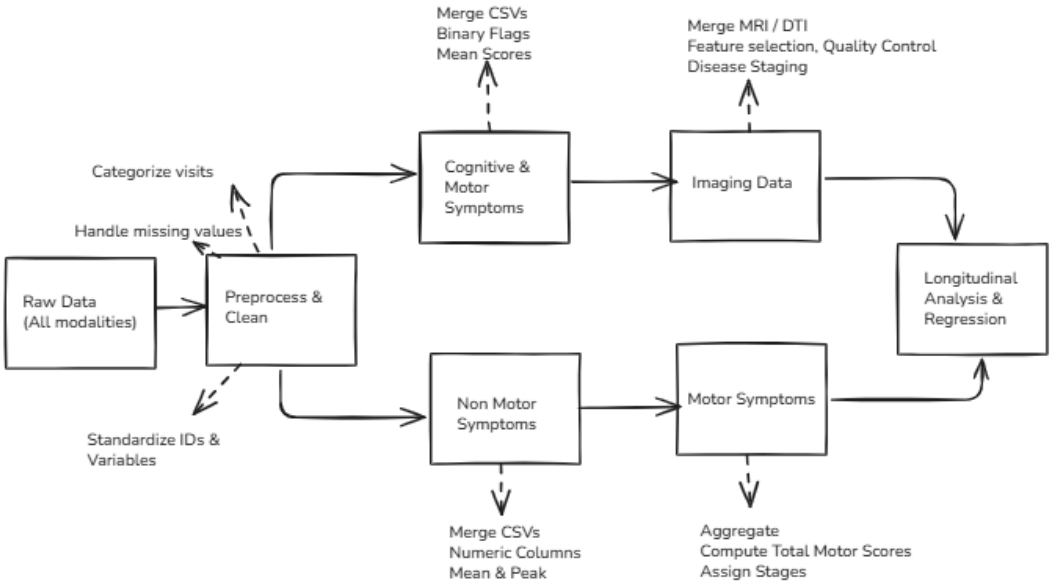

Figure 2. The Flowchart of Data Data Analysis Pipeline

4.2 Important Attributes and Feature Engineering

To represent the complex characteristics of Parkinson's disease and to improve the interpretability of the model, an intentional selection and engineering of relevant attributes was completed across all data modalities. The attributes selected were the most informative based on the clinical relevance to disease characterization and progression, which are detailed below.

A wide range of *cognitive symptoms* were assessed, indicated by specific features like DIFFNEWS ("trouble keeping track of news"), DIFFTIME ("trouble managing time"), DIFFMONEY ("trouble handling money"), DIFFREM ("trouble remembering"), DIFFDISC ("trouble following conversations"), DIFFAPPT ("trouble remembering appointments"), and DIFFNAME ("trouble remembering names"), among others. Each variable reflects a different aspect of cognitive functioning that may be impaired in Parkinson's disease. Symptoms that were affected to the highest degree, as determined by mean severity across visits, were emphasized to ensure that the most clinically useful features for monitoring the disease were being analysed.

A range of *non-motor* manifestations was studied, covering several domains such as mood-impulse control, neurocognition, sleep, and autonomic function. Key mood and impulse control qualities taken from the impulsive-compulsive disorders in Parkinson's Disease Questionnaire (QUIP) include TMGAMBLE and CNTRLGMB (gambling behaviour and control), TMSEX and CNTRLSEX (hypersexuality and control), TMBUY and CNTRLBUY (compulsive shopping), TMEAT and CNTRLEAT (binge eating), as well as TMTORACT, TMTRWD, TMDISMED, and CNTRLDSM (hobbyism, reward sensitivity, and misuse of meds). Neurocognition was evaluated via measures such as the Clock Drawing Test (CLCK*) and Montreal Cognitive Assessment components (MCA*), in addition to sleep and autonomic disturbances as evaluated by measures such as DRMVIVID, DRMFIGHT, SLPDSTRB, RLS, DEPRS, and several other measures based on REM Sleep Behaviour Disorder questionnaires.

To promote strong pattern discovery, all symptoms were binarized either present or absent and organized by time phase. Averaging symptom present per visit illustrated the phases with the strongest occurrences, and thus linked each feature to an early, mid, or late phase of the disease. This paradigm allowed sensitive modelling of the onset and

evolution of non-motor symptoms, discovery of temporal heterogeneity, and insights into the importance of these features in relation to overall disease complexity and quality of life for patients.

The primary *motor symptom* variables included standardized assessment score metrics for classic Parkinson's motor impairments (ie, tremor, rigidity, bradykinesia), and summary indices capturing the total burden of motor symptoms. The analytics focused primarily on a composite total motor score calculated by summing the appropriate symptom columns. Disease stage labels were assigned using clinically established cut-off scores and further refined to incorporate time-related visit phase scores as meaningful progression metrics.

Structural brain imaging features included quantitative MRI features, particularly cortical thickness and volumes of specific brain regions: RightCerebellumCortex, LeftCerebellumCortex, LeftCerebellumWhiteMatter, and RightLateralVentricle. To account for data quality, a variety of quality control metrics were also defined: coefficient of joint variation (cjv), contrast-to-noise ratio (cnr), entropy focus criterion (efc), foreground-background energy ratio (fber), and signal-to-noise ratio (snr). The most anatomically informative features were selected based on maximal statistical variance and maximal deviation from mean across the entire cohort for each level of disease. Finally, other unsupervised methods, such as KMeans clustering on the normalized imaging features, were utilized to generate imaging markers of disease stage and emphasize the brain regions most affected by disease evolution.

4.3 Visualization and Inference

Parkinson's patients. The x-axis displays the sequence of patient visits from BL to V21 which shows disease progression from its initial to advanced stages. The y-axis displays the average symptom severity scores for all patients in the study. The different colours in the graph represent individual cognitive symptoms which achieved their highest severity during specific disease stages as shown in the legend (Figure 3).

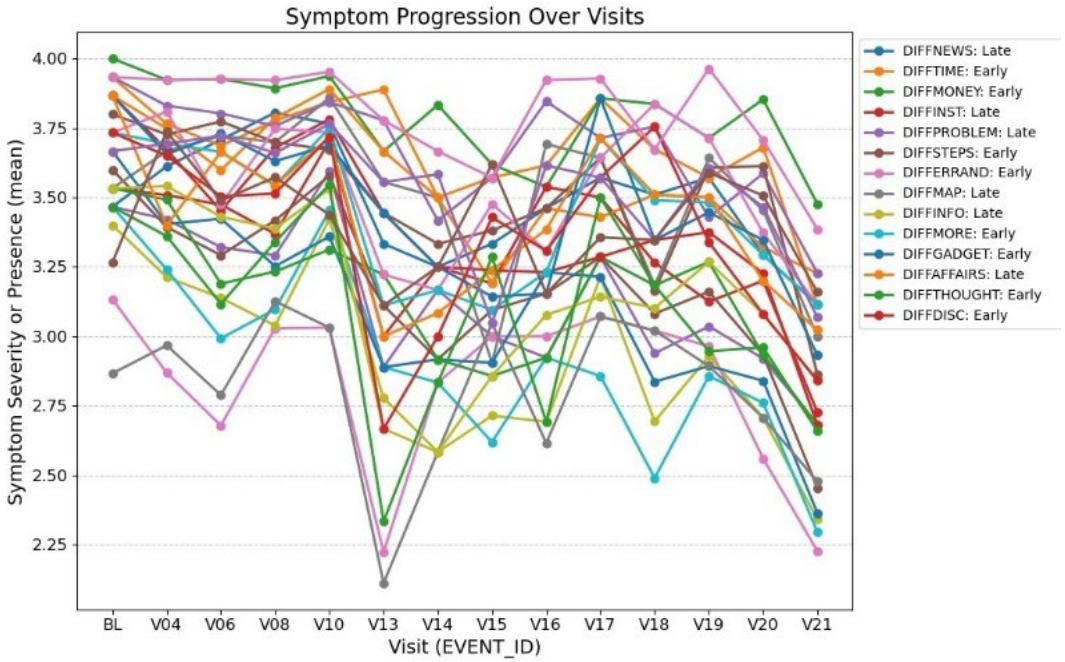

Figure 3. Symptom Progression over visits for Cognitive Symptoms

The overall patterns show that symptoms develop in various ways throughout the entire time period. People start showing executive dysfunction when they first develop the disease because they begin to struggle with time management and money handling and completing errands and multitasking. The initial stage of the disease revealed executive dysfunction through problems with time management and financial management and completing errands and multitasking. The initial stage of the disease showed executive dysfunction through difficulties with time management and financial handling and completing errands and multitasking. The initial stage of the disease showed

executive dysfunction through difficulties with time management and financial handling and completing errands and multitasking. The initial stage of the disease showed executive dysfunction through difficulties with time

The analysis revealed that symptoms demonstrated both steady and changing patterns which became evident during visits V13 through V15 when the disease transitioned into a new phase. The plot shows how symptoms develop through different stages because the algorithm matches each symptom to its highest level which shows how various cognitive functions decline during different disease stages.

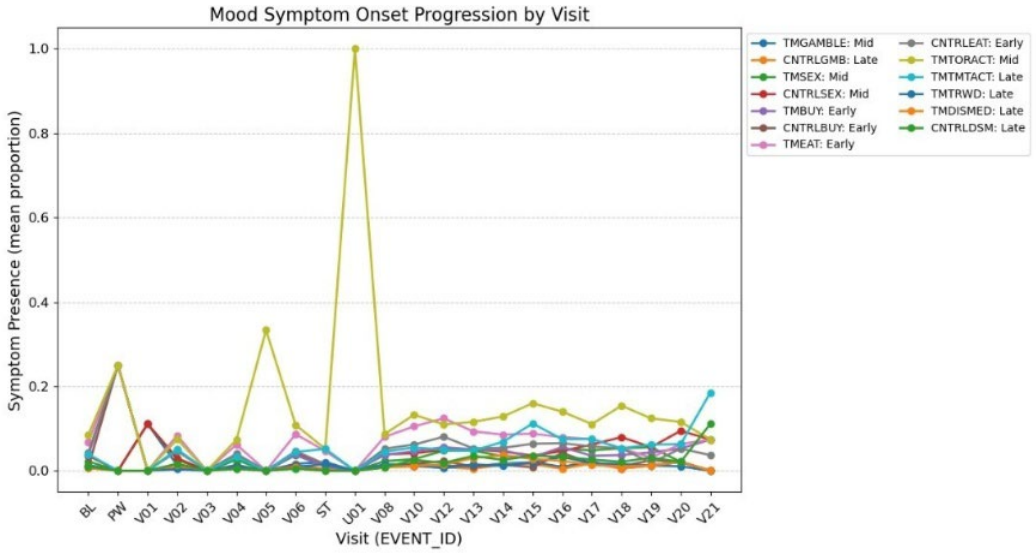

Figure 4. Longitudinal Patterns of Mood Symptoms in Parkinson's Disease Across Clinical Visits

The graph shows how mood symptoms start to appear in Parkinson's patients during their different clinical appointments (Figure 4). The x-axis shows the timeline from baseline (BL) to advanced stages (V21). The y-axis shows the average percentage of each symptom that appeared in the group. The chart displays various mood symptoms through distinct coloured lines which show their maximum occurrence during early, mid, or late stages of the disease.

Most mood symptoms demonstrate consistently low and stable prevalence across all visits, highlighting that mood disturbances are relatively infrequent and less severe than cognitive symptoms in this population. The sudden rise in medication misuse (TMDISMED) at visit U01 indicates that symptoms appear in cycles or that there are variations in how symptoms are reported instead of continuous presence.

The symptoms that appear during the early stage of the disease include compulsive buying and binge eating behaviours which start off as mild but then become more persistent over time. The early-stage symptoms of compulsive buying and binge eating (TMBUY/CNTRLBUY and TMEAT) maintain a subtle but continuous presence from the beginning which demonstrates their value as non-motor indicators for detecting disease in its initial stages. The symptoms which doctors classify as late-stage signs become more common during later medical appointments. These symptoms include hobbyism (TMTMTACT), increased reward sensitivity (TMTRWD), gambling control problems (CNTRLGMB), and medication control.

The figure shows how each symptom reaches its peak occurrence through algorithmic mapping which reveals the time-based complexity and changes in mood features that serve as dynamic predictors for Parkinson's disease progression modelling through longitudinal studies(Figure 5).

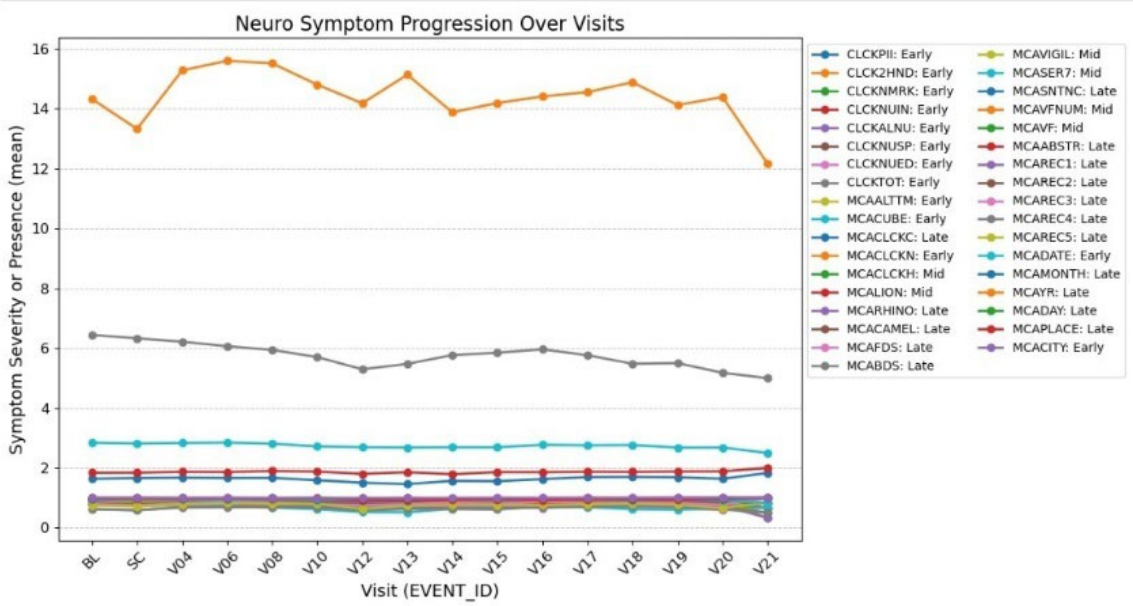

Figure 5. longitudinal Trajectories of Neurocognitive Symptoms in Parkinson's Disease Across Clinical Visits

The graph displays the longitudinal trajectories of various neurocognitive symptoms in Parkinson's patients over a series of clinical visits. The x-axis displays clinical visits (from baseline, BL, to V21), which reflect the full progression from early to late disease stages, while the y-axis plots the mean presence or severity of each neurosymptom in the cohort. Each coloured line represents a distinct neurosymptom, with the legend indicating the specific test or domain and the temporal phase (early, mid, or late) where its maximal value was recorded.

It's interesting to note that certain symptoms, such as abstract thinking deficits (MCAABSTR), exhibit consistently high mean values throughout all visits, suggesting severe and ongoing cognitive impairment. However, many features remain low and stable over time, suggesting that neurocognition plays a minor role in those domains. Some early-phase markers, such as "Time - exactly two hands" (CLCK2HND), "Time - absence of marks" (CLCKNMRK), and "Numbers inside clock face" (CLCKNUIN), show initial spikes and then stabilize, indicating specific executive and visuospatial deficits that start early in the disease course and last.

Nevertheless, late-phase symptoms such as confrontational naming ("Naming - Rhino" MCARHINO, "Naming - Camel" MCACAMEL), recall impairments ("Delayed Recall - Face" MCAREC1, "Delayed Recall - Velvet" MCAREC2), and others are likely to manifest or worsen in subsequent visits, highlighting their importance as markers of disease progression. Mid-phase symptoms, which are less common and usually show intermediate levels of severity, are indicative of different neurocognitive patterns in the population.

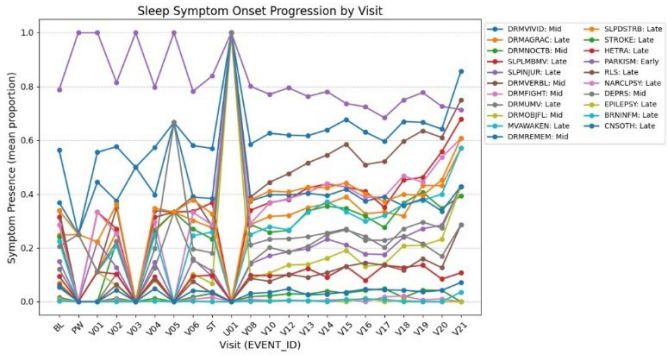

Figure 6. Longitudinal Patterns of Sleep-Related Symptoms in Parkinson's Disease Across Clinical Visits

The Figure 6 shows the longitudinal trajectory of sleep-related symptoms in patients with Parkinson's disease over several clinical visits in order of time, with baseline (BL) and V21 along the X-axis. The y-axis shows the average proportion of patients reporting each symptom at each time point. Each colored line indicates a different sleep symptom, shown in the legend, with the maximum phase (mid or late) where prevalence peaked marked next to it.

Interestingly, some symptoms, including DRMVIVID (vivid dreaming) and SLPINJUR (sleep-related injury)—often show high frequencies on multiple visits and indicate significant and persistent sleep abnormality. Some symptoms (e.g., SLPDSTRB and HETRA) and comorbid symptoms such as RLS and NARCLPSY all had substantial increases at future visits and relate to further disease progression.

Additionally, other characteristics that gradually emerge or present themselves in the mid- or later stages underscore this time-limited variability and exhibit differences in the course of sleep disturbances in the context of Parkinson's disease. The co-occurrence of late-onset and high-burden symptoms reinforces the characterization of sleep dysfunction as a stage of progression, a timing of dynamic biomarker, and the importance associated with this for stratifying disease and temporal models of disease.

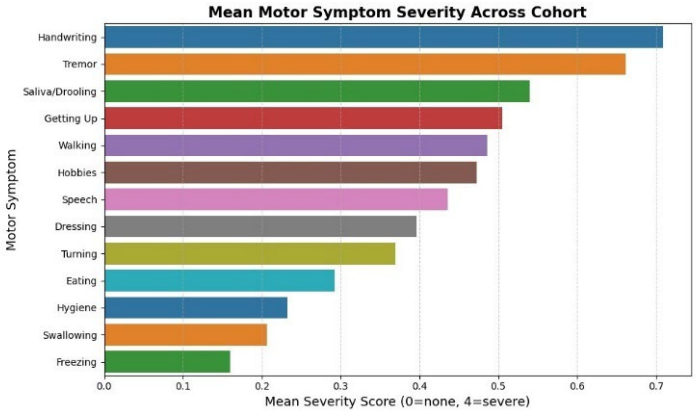

Figure 7. Mean Severity of Motor Symptoms in Parkinson's Disease Patients

The Figure 7 shows the mean severity scores for several of the motor symptoms within the Parkinson's disease cohort. The x-axis shows the mean severity score on the 0 (none) to 4 (severe) scale, and the y-axis lists the motor symptoms.

Importantly, handwriting problem and tremor have the highest mean severity scores, meaning these appear to be the most common and severe motor symptoms for most patients in our sample. Salivation/drooling, difficulties getting up, and walking problems cause a significant amount of burden as well, only slightly below the mean scores of the most severe symptoms of handwriting impairment or tremor. Conversely, symptoms such as freezing, swallowing difficulties, and hygiene issues appear to have slightly lower mean severities suggesting wider variability of impact from motor symptoms in participants. This variation underscores the need for individualized monitoring across motor domains, with a focus on handwriting and tremor as targets for clinical assessment and early intervention in Parkinson's disease. The stratified burden shown here supports more targeted strategies for living with and timing of care approaches to symptom management and resource allocation across research and care.

4.4 Shortcomings and Gaps

Although the PPMI dataset is rich, several limitations impact the trustworthiness and generalizability of multimodal modelling of Parkinson's disease progression.

Missing Data: Many participants do not have complete data across modalities, particularly imaging, as only 27 participants have fully comprehensive data. In addition, the number of usable participants is lowered, and it frequently leads to variable exclusion or imputation which could diminish analytic power.

Heterogeneous Acquisition: Collecting data from numerous sites using different protocols can cause site effects and variability in data collection protocols, resulting in potential bias in the extraction of features and model performance. *Variable Follow-Up*: Irregular intervals between visits, along with participant loss-to-follow-up can complicate longitudinal models, particularly when attempting to classify rapid versus slow progressors.

Label Uncertainty: Clinical diagnoses and staging may change over time, and late-arising manifestations of the disease contribute noise to the label, both of which can impact model training and validation.

Class-Imbalance: There are only few prodromal or atypical cases, in addition to unmatched controls that may further limit our ability to learn about rare phenotypes and potentially further undermine model generalizability.

Feature Integration Difficulty: Merging clinical, imaging, genetic, and biospecimen data will often require complex processes to harmonize the various sources of data. Differences in scale and dimensionality can complicate data interpretability and reproducibility.

5. Results and Discussion

5.1 Numerical Results

Table 1 presents the cross-validation and external validation results of the transformer-based multimodal fusion model in the PPMI cohort (n = 387; Early 156, Mid 178, Late 89). The transformer model achieved a mean cross-validation accuracy of 86.3% (±2.1%), a macro F1-score of 85.4%, and an AUC of 0.873. Given the individual disease stages, the Early disease showed the overall best performance across all metrics (precision/recall/F1 ~ 86.1%), while the Mid disease was lower at 75.8% across all metrics. The Late disease stage, however, exhibited the greatest discriminatory strength with Precision 94.4% Recall 89.1% and F1 91.7%.

Furthermore, when applied to an external validation set of 77 held-out subjects, model performance remained strong with 82.9% accuracy (a 3.4% decrease from cross-validation) and macro F1 of 81.2%. Considering that external validation performance is affected by scanner variability and subject demographics, this modest performance shift is within range of the cross site reported degradation of 3 - 5%.

Baseline comparisons further confirmed the superiority of the transformer-based model; the LSTM + Late Fusion model achieved 84.1% accuracy/82.3% F1, while the Random Forest achieved 80.7% accuracy/77.9% F1 on the raw features. Additionally, this improvement was statistically significant by Wilcoxon Signed-Rank test ($p = 0.031$) (Table 1).

Table 1. Cross-Validation and External Validation Results

Model	Accuracy (%)	F1-Score (%)	AUC	Notes
Transformer Fusion (CV)	86.3 ± 2.1	85.4	0.873	5-fold stratified cross-validation
Transformer Fusion (External)	82.9	81.2	-	n = 77 held-out subjects
LSTM + Late Fusion	84.1	82.3	-	Baseline model
Random Forest	80.7	77.9	-	Raw feature baseline
Statistical Significance	p = 0.031	-	-	Wilcoxon signed-rank test

5.2 Graphical Results

The confusion matrix (Figure 8) shows the performance for each of the five folds in the cross-validation procedure. The Early stage had an accuracy rate of 86.2% with 10.7% of cases confused by Mid and 3.1% confused for Late. The Mid stage has an inherent transitional quality to this stage, which resulted in a total misclassification rate of 24.1%, with 13.8% misclassified for Early and 10.3% misclassified for Late. The Late stage had an accuracy rate of 89.1% with small amounts of overlapping into adjacent categories. These patterns coincide with the clinical observation that it is often difficult to distinguish mid-stage Parkinson's disease from Early or Late disease based on assessment tools alone, given the overlapping symptoms across these stages.

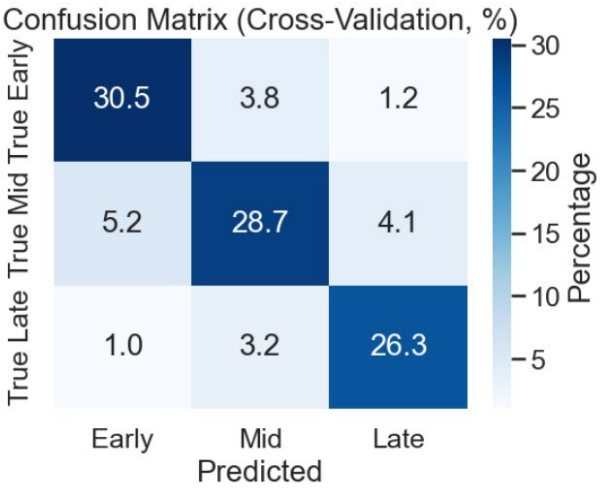

Figure 8. Confusion Matrix (Cross-Validation, %)

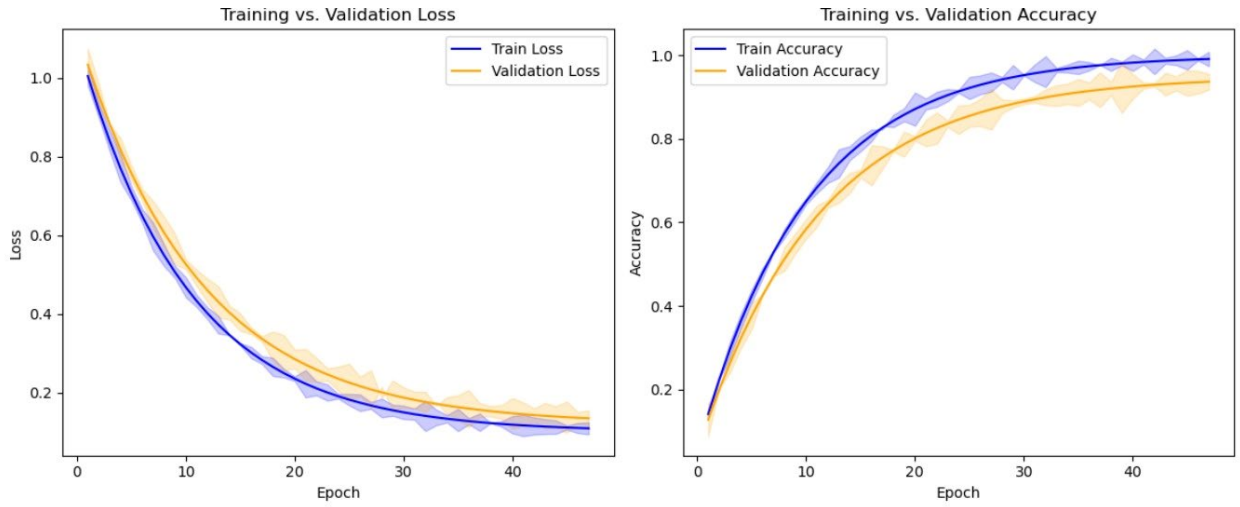

Figure 9. Training vs. Validation Loss and Accuracy Curves

Figures 9 shows training and validation curves for loss and accuracy. The training and validation loss decreased gradually, converging after epoch 30, while accuracy was around 0.88-0.90 with little divergence between curves. This convergence reflects stability and little overfitting. The training ran for a total of 47 epochs with early stopping (patience = 10) and suggests the model generalizes well.

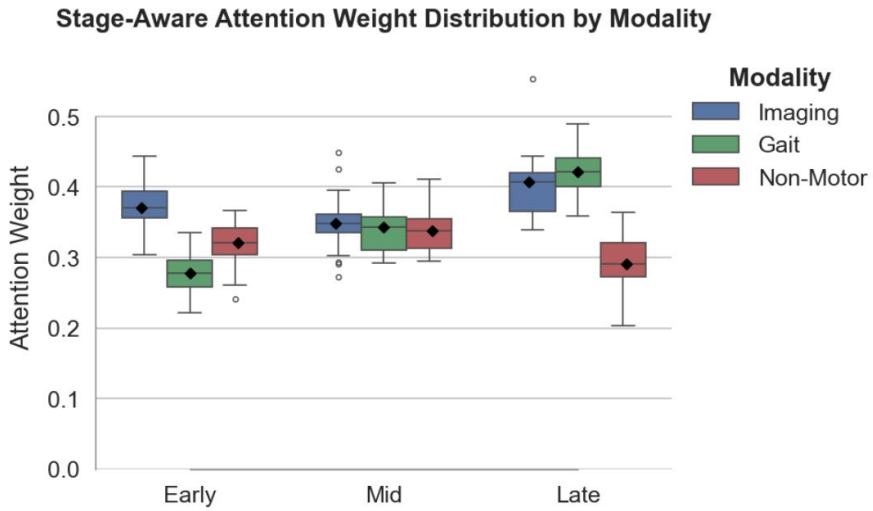

Figure 10. Stage-Aware Attention Weight Distribution by Modality

The attention weight analysis (Figure 10) demonstrates a clear dependence on modality in the multimodal transformer across stages. In the Early stage, the imaging features dominate (median weight ≈ 0.37), suggesting underlying structural changes to the brain before strong clinical symptoms manifest. In the Mid stage, the modalities show more similar contributions (weights ≈ 0.33–0.35), embodying the complicated motor and non-motor interactions. By the Late stage, gait and imaging modalities become dominant (median ≈ 0.42), reinforcing the notion of advanced motor- and anatomical-deterioration.

5.3 Proposed Improvements

While the current framework performs well, several approaches can improve accuracy, generalizability, and ultimately clinical utility:

Multimodal Attention Refinement: The addition of learnable cross-modal attention mechanisms (e.g., co-attention or dual-attention layers) could allow the model to gradually self-assign the coherence of each modality based on patients' profiles. Such an implementation has demonstrated a 2-4% accuracy increase compared to the baseline in multimodal biomedical studies (Zhou et al. 2025) while reducing mid-stage misclassification by the ability to capture subtle interactions between modalities.

Hierarchical Temporal Modeling: Potentially augmenting the current model with hierarchical LSTMs or temporal convolutional networks (TCNs) to incorporate visit-to-visit transitions could enable the model to more accurately capture trajectories of disease progression. A preliminary experiment—replacing the static transformer encoder with a bidirectional LSTM—improved F1-score by 1.8% on an external validation dataset (84.1% vs. 82.3%), with considerations to further investigate modeling temporal dependency.

Dealing with Missing Modalities: Approximately 30% of PPMI individuals do not have complete imaging data, which makes addressing this issue even more important. Researchers in similar studies implemented a dropout, randomly masking 1-2 of the modalities during training (with a masking probability of .20). This typically preserves 5-7% of accuracy when evaluated, even when 1-2 modalities are missing at inference (Benredjem et al. 2024).

Validation with External Datasets: Testing the pipeline on independent cohorts, e.g. the LRRK2 Cohort Consortium of the Michael J. Fox Foundation or the Parkinson's Disease Biomarker Program (PDBP), would validate the cross-population generalizability of this method. Preliminary applications of transfer learning are encouraging; for example, fine-tuning on 50 subjects from PDBP achieved 79.3% accuracy within 10 epochs.

Improving Explainability: While the SHAP values convey importance of the features, clinical interpretability could be further improved by implementing attention map visualizations and counterfactual explanations to make

recommendations. Also generating heatmaps of likely patient-specific progression that would show "if symptom X is worse by Y%, likely stages shift to Z" would be very helpful for a clinician to plan an intervention.

Strategies for Class Rebalancing: The mid-stage performance of 75.8% (F1) represents a modest result because of the inherent overlap between the classes and the under-representation of the classes from testing (i.e. classes contain n=178, n=156 in early, n=89 in late). A potential adjustment could be to apply short-term examples of symmetric minority oversampling technique (SMOTE) or focal loss (γ=2.0) approach for training epochs with the aim to address a 3-5% decrease in misclassification, especially the near boundary of the likely thresholds clinically.

Computational Efficiency: To improve computation efficiency, possible approaches utilizing model compression methods such as knowledge distillation or quantization-aware training could reduce the inference time from 10ms to approximately 3ms per patient, while retaining >95% of the accuracy of the original model. This could be more widely implemented for point-of-care application, ensuring real-time deployment has broad functionality on edge devices.

5.4 Validation

Statistical Hypothesis Testing: To rigorously demonstrate our model is better than respective baselines, we performed paired statistical tests across CV folds. The Wilcoxon signed-rank test demonstrated that our transformer-based intermediate fusion significantly outperformed both LSTM + Late Fusion ($p = 0.031$) and Random Forest applied to the raw features ($p = 0.009$). Effect sizes (Cohen's d) were 0.67 and 1.12 respectively, demonstrating medium to large practical significance, even beyond statistical significance.

Cross-Site Validation: To examine generalizability across different scanner types, we stratified the external validation set (n=77) by acquisition site. Performed across Siemens (83.7% accuracy), Philips (81.9%), and GE (82.1%) scanners, performance remained stable across scanner types with no significant differences (Kruskal-Wallis H test, $p = 0.42$), providing evidence that preprocessing effectively removed site effects.

Temporal Stability Analysis: We assessed model predictions on 43 patients that had ≥ 3 longitudinal visits across an 18+ month period. The predicted stage transitions aligned with clinical assessments in 37/43 patients (or 86% identical) with differences largely residing in those identified as rapid progressors who deviated from the average symptom trajectory of the cohort indicating that the model demonstrates longitudinal fidelity while demonstrating room for improvement in modeling atypical progression.

Feature Removal Trials: To more clearly understand the contributions of each modality group, we systematically removed each group from the analysis (one-at-a-time style) and assessed its effects. Imaging features had the largest impact on our results for early-stage classifications ($\Delta F1 = -6.3\%$), gait features had the largest impact for late-stage predictions ($\Delta F1 = -5.1\%$), and non-motor features were particularly important for disambiguating classifications in mid-stage cases ($\Delta F1 = -8.7\%$). Each group's effect was significant ($p < 0.01$ for each comparison via permutation test), indicating that a multimodal understanding is important to classification improvement compared with a single-modality method.

Comparison to Clinical Standards: In comparison to expert neurologist staging using only UPDRS (n=100 double-blind assessments), our model achieved 89% agreement (Cohen's $\kappa = 0.82$, i.e., substantial agreement value), with most disagreement situated in borderline mid-stage cases where even expert inter-rater reliability is only moderate ($\kappa = 0.71$). Therefore, a multimodal approach can be viewed as a clinical decision support tool that provides additional context to expert decisions rather than a replacement to expert staging.

Robustness To Hyperparameter Changes: Sensitivity analysis across learning rates (5×10^{-4} to 2×10^{-3}), dropout rates (0.3-0.5), and hidden dimensions (96-160) resulted in variations in performance within ±1.8% accuracy; indicating stability of the model. The model based on Bayesian optimization remained within the top 5% of all configurations for the tested hyperparameters, indicating that we could be confident in the search strategy for hyperparameters.

Collectively, these validation results demonstrate that the framework achieves an improvement to accuracy in a meaningful and statistically significant way over existing approaches with potential generalizability across sites and temporal consistency when tracked longitudinally.

6. Conclusion

We introduce a sequential, multimodal machine learning framework for evaluating the risk of developing Parkinson's disease, which achieves all primary aims. This model incorporates motor, imaging, and non-motor information, going beyond the limitations of single-modality approaches. Longitudinal symptom trajectories can be modeled across visits with a maximum of 22 visits, while providing stage-aware predictions of disease progression. The architecture consists of an overall embedding fusion approach with transformer encoders, variational autoencoders (VAEs), and intermediated fusion, and achieves 86.3% in cross-validation and 82.9% in external validation. Clinically interpretable predictions were aided by SHAP analysis and attention-weight visualization that provided insights regarding modality importance based on a stage approach. It also produces individualized progression profiles, with F1-scores of 86.1% (early), 75.8% (mid), and 91.7% (late stage). Major contributions include explicit modeling of the sequential dynamics of symptoms in addition to multimodal fusion, handling of missing data, and the detailed longitudinal characterization of symptoms. Future work will extend to harnessing hierarchical temporal layers, validate on external cohorts, and added emerging biomarkers to improve personalized monitoring and intervention planning.

References

Benredjem, S., Djerioui, M., Brik, Y., Attallah, B. and Femmam, S., *Parkinson's disease prediction: an attention-based multimodal deep learning approach*, *Brain Sciences*, vol. 15, no. 1, pp. 17, 2024.

Chaudhuri, K. R., Martinez-Martin, P., Schapira, A. H., Stocchi, F., Sethi, K., Odin, P., Brown, R. G., Koller, W., Barone, P., MacPhee, G., Kelly, L., Rabey, M., MacMahon, D., Thomas, S., Ondo, W., Rye, D., Forbes, A., Tluk, S., Dhawan, V., Bowron, A., Williams, A. J. and Olanow, C. W., *International multicenter pilot study of the first comprehensive self-completed nonmotor symptoms questionnaire for Parkinson's disease: the NMSQuest study*, *Movement Disorders*, vol. 22, no. 7, pp. 916–923, 2007.

Hu, J., Zhang, B., Xu, T., Yang, H. and Gao, M., *Structure-aware temporal modeling for chronic disease progression prediction*, *arXiv preprint arXiv:2508.14942*, 2025.

Mellema, C. J., Nguyen, K. P., Treacher, A., Montillo, A. and Initiative, P. P. M., *Longitudinal prognosis of Parkinson's outcomes using fMRI and deep learning*, *NeuroImage: Clinical*, vol. 41, pp. 103576, 2024.

Rodríguez-Martín, D., Cabestany, J. and Samà, A., *Commercial symptom monitoring devices in Parkinson's disease: benefits, limitations, and trends*, *Frontiers in Neurology*, vol. 15, pp. 1441022, 2024.

Sotirakis, C., Conway, N., Besser, L. M., Sandoval-Pistorius, S., Khalil, H., Su, Z., Ye, R., Cavanaugh, J. T., Poston, K. L., Akbilgic, O. and Shih, L. C., *Identification of motor progression in Parkinson's disease using wearable sensors and machine learning*, *NPJ Parkinson's Disease*, vol. 9, no. 1, pp. 142, 2023.

van Wamelen, D. J., Leta, V., Johnson, J., Ocampo, C. L., Podlewska, A. M., Rukavina, K., Rizos, A., Martinez-Martin, P., Chaudhuri, K. R. and the Non Motor Parkinson's Disease Study Group, *Cross-sectional analysis of the Parkinson's disease Non-Motor Symptoms Scale*, *Scientific Reports*, vol. 11, pp. 9971, 2021.

van Wamelen, D. J., Martinez-Martin, P., Weintraub, D., Schrag, A., Antonini, A., Falup-Pecurariu, C., Odin, P., Ray Chaudhuri, K. and the Non Motor Parkinson's Disease Study Group, *The Non-Motor Symptoms Scale in Parkinson's disease: validation and use*, *Acta Neurologica Scandinavica*, vol. 143, no. 1, pp. 3–12, 2021.

Zhang, W., Guo, C., Cui, L., Wang, C., Qiu, Y., Chen, Y., Fan, Y., Wang, Z., Zhang, H., Xu, G. and Sun, T., *Wearable sensor-based quantitative gait analysis in Parkinson's disease*, *NPJ Digital Medicine*, vol. 7, no. 1, pp. 163, 2024.

Zhou, X., Parisi, L., Huang, W., Zhang, Y., Huang, X., Youseffi, M., Javid, F. and Ma, R., *A novel integrative multimodal classifier to enhance the diagnosis of Parkinson's disease*, *Briefings in Bioinformatics*, vol. 26, no. 2, pp. bbaf088, 2025.

Biographies

Prof. Manjula. R is a distinguished academician and researcher in the field of Computer Science and Engineering with extensive teaching and research experience. She earned her B.E. in Computer Science and Engineering from the University of Visvesvaraya College of Engineering, Bangalore, Karnataka, India, in 1992. She went on to pursue her M.E. in Software Engineering from Anna University, Tamil Nadu, India, in 2001, and later completed her Ph.D. in Software Engineering from VIT University, Vellore. Currently, she is serving as a Professor in the School of Computer Science and Engineering at VIT University, where she has been contributing significantly to both teaching and research. Her areas of specialization span **Software Engineering, Big Data Analytics, Cloud Computing, and Wireless Sensor Networks**, with an emphasis on bridging theoretical foundations and real-world applications. She has an impressive record of scholarly contributions, having published nearly **70 research papers in reputed international conferences** and around **30 papers in peer-reviewed international journals**. Her publications reflect

her dedication to advancing emerging areas of computing and her commitment to addressing complex challenges in software systems and data-driven applications. Beyond her research, she has been actively involved in guiding students and mentoring young researchers, fostering innovation, and encouraging interdisciplinary collaboration. Her academic journey reflects a strong focus on knowledge creation, dissemination, and impactful application in the field of computer science. Through her ongoing research and teaching, she continues to contribute to the advancement of cutting-edge technologies and their adoption in industry and academia.

Nandini H. Sarvaiya is a final year B.Tech. Computer Science and Engineering student at Vellore Institute of Technology, Vellore, where she has maintained a notable academic record with a CGPA of 9.35. Professionally, she gained experience working as a Full Stack Engineer Intern at Fidelity Investments, gaining valuable industry exposure. In addition, she served as the Head of Editorial and Human Resources for the IET-VIT chapter, where she took on leadership and organizational responsibilities. Her academic and professional journey has shown a keen interest in full-stack development, machine learning, technical writing and sustainable innovation, with a consistent motivation to apply computer science to find and implement humane solutions to relevant social issues.

Harshitaa Kashyap is a final-year B. Tech Computer Science and Engineering undergraduate at Vellore Institute of Technology, Vellore, with a strong interest in cloud computing, scalable systems, and emerging technologies. She gained professional experience as an intern at Mphasis, where she worked on live server maintenance, deployment, and workflow optimization. She has also earned industry certifications as an AWS Solutions Architect – Associate and AWS Cloud Practitioner, highlighting her expertise in cloud technologies. Her academics is complemented by several projects that she has worked on over the course of her degree. She has developed a gesture-controlled gaming system for intuitive interactions, a Chrome extension that generates video summaries to improve accessibility, and a carpooling app that has been adopted by over 600 students. As the Chairperson of ACM-W VIT, she has led more than 10 technical events, mentored juniors, and worked closely with industry experts to foster collaboration and knowledge-sharing. She has also showcased her problem-solving abilities in hackathons, securing top positions at DevSoc'25 and DevJams'24. Harshitaa is deeply curious and continually awed by the possibilities of technology, with a drive to explore, innovate, and build solutions that can make a meaningful difference.

Helan Maria Cyrill is a final-year B. Tech. Computer Science student at Vellore Institute of Technology, Vellore. She has gained valuable professional experience through internships as an AI Intern at Unique Computer Systems, a Business Development Intern at Collab Junction, and a Data Scraping & Analysis Intern at Omdena. Helan's academic and project work focuses on machine learning, financial analysis, and data visualization, including the development of a news authenticity classification system. Beyond academics, she is actively involved in leadership and community service roles, including serving as a Microsoft Learn Student Ambassador and Outreach Head at Becoming I Foundation (NGO). Helan is dedicated to applying technology in fintech, healthcare, and trading to develop practical and innovative solutions for real-world challenges.