A Lightweight Self-Supervised Contrastive-Inspired Framework for Interpretable Parkinson's Disease Staging using Wearable Gait Signals

Gunjan Mittal Roy[1], Ajay Kumar[2]
[1]Lincoln University College Malaysia, [2]IILM University, Greater Noida, India
[1]Email ID: gunjanmittal2013@gmail.com; [2]Email ID: ajay.phdcse@gmail.com

Abstract

Parkinson's disease (PD) is characterized by progressive neurodegenerative changes resulting in movement impairments, instability during gait, shaking, and problems with posture. Accurate staging of the disease is necessary for timely and efficacious treatment interventions; however, current standardized clinical assessments are often subjective, time-consuming and have considerable inter-rater variability. Although the integration of wearable sensing technologies with AI for automated gait analysis has become possible, these approaches typically require supervised learning, large training sets of labelled data and have limited interpretability in their models.

This study presents a lightweight, self-supervised, contrastive-inspired approach to interpreting PD staging through wearable gait data. The framework incorporates data augmentation and a lightweight multi-layer perceptron (MLP) as an encoder to produce latent representations of augmented gait data and applies classical ML classifiers in combination with explainability via SHAP to a unified architecture. By producing discriminative latent embeddings from augmented gait data that do not require extensive amounts of labelled data, our framework effectively captures comorbidity associated with disease progression through gait data. These latent embeddings are classified using Random Forest (RF) and Support Vector Machine (SVM) classifiers, with RF yielding superior classification performance within the context of our proposed framework.

Results from the experimental study demonstrate RF and SVM classifiers using our framework achieved ~95% accuracy across all three classes while outperforming the baseline ML classifier and being more robust to imbalanced class distribution. The PCA biplots indicate that our framework produced clearly defined embeddings between the three classes, while the confusion matrix indicates that the multi-stage classification results were accurate. Additionally, SHAP-based explainability identified clinically relevant gait features that were related to the progression of the disease and thereby enhanced the transparency of our framework, leading to improved trust in clinical decision-making.

The current framework provides an efficient, interpretable, and reliable system for stage classification of Parkinson's disease, thus it can be used as a solution to monitor disease progression in real time and to allow for wearable health monitoring.

Keywords: Parkinson's disease, wearable gait analysis, self-supervised learning, contrastive learning, latent embeddings, Random Forest, SHAP, explainable artificial intelligence (XAI).

1. Introduction

Millions of people around the world suffer from Parkinson's disease (PD), which is one of the most common neurodegenerative diseases. It primarily affects movement, as PD is caused by the destruction of dopamine-secreting neurons in the substantia nigra region of the brain. PD produces many different symptoms, including tremor, bradykinesia, rigidity, poor posture, and gait problems. Gait abnormalities have been recognised as being among the first and most reliable signs that PD is progressing [1], [2]. Clinically diagnosing and staging PD is generally done using standardised assessment measures such as the Unified Parkinson's Disease Rating Scale (UPDRS) and the Hoehn and Yahr stage [3]. Unfortunately, both of these methods use subjective judgement from the clinician performing the examinations and are only completed at a set time period, which limits scalability and reliability of these assessments. As a result, continuous monitoring of gait using wearable sensors has emerged as an effective means of objectively measuring signs of PD [4].

Recent advances in machine learning and deep learning methods have significantly improved the ability to use gait signals to automatically diagnose PD. SVM, RF, and CNNs have been used by researchers to classify PD in many studies [5]–[7]. Although the research supporting these classifications has shown great promise, much of it has focused primarily on binary classifications rather than multi-stage analyses. Additionally, deep learning classifiers generally require large labelled training datasets, and they are poorly interpretable, which limits their usability in healthcare settings.

Unsupervised methods, such as clustering, have been attempted to find latent stages of progression [8, 9]. The output of any clustering method is highly contingent upon the effectiveness of the handcrafted measurements or features, and the algorithms used for clustering cannot learn to represent gait dynamics discriminatively. As a result, their ability to generalize will be very limited. Recently, self-supervised learning has been identified as a strong approach to developing new representations since self-supervised learning does not require large amounts of labeled data to learn representations. Contrastive methods, such as SimCLR, MoCo, and Bootstrap Your Own Latent (BYOL), have been shown to perform well in computer vision and signal processing applications [10-12]. Contrastive methods represent data using latent representations created by comparing the similarity of multiple augmented copies of the same instance against the similarity of augmented copies that are from different instances.

Despite these advances in developing latent representations using contrastive learning methods, there has been limited research on using lightweight self-supervised representation learning methods to create models for Parkinson's disease classification using gait data. Additionally, the incorporation of explainability in self-supervised model development is a critical aspect of establishing credibility and use within the clinical environment. This paper addresses these two gaps in the development of interpretable models for classifying Parkinson's disease from wearable-based gait data by presenting a lightweight, self-supervised contrastive-inspired model for the classification of Parkinson's disease based on valuable accelerometer-based data obtained from wearables. The proposed framework consists of Data augmentation, lightweight representation learning, classical machine learning classifiers, and Explainable AI Techniques.

Novelty and Major Contributions of the Proposed Work

The originality of the very innovative work lies in creating a novel lightweight self-supervised contrastive-inspired framework to enable interpretable staging of Parkinson's disease using wearable gait signals. The proposed method differs significantly from conventional supervised methods and clustering-based methods in that the author's proposed method integrates representation learning, classical machine learning, and explainable AI methods into a unified computationally efficient architecture that is specifically tailored for addressing the common healthcare challenges of limited labelled data, poor interpretability, and class imbalance.

The key contributions of the very innovative work are as follows:

1. A novel lightweight self-supervised contrastive-inspired embedding framework for discriminative gait representation learning with reduced reliance upon extensive labelled datasets has been developed to improve feature learning and generalization performance.
2. Random Forest and Support Vector Machine classifiers have been integrated to accurately and robustly stage Parkinson’s disease using learned latent embeddings derived from wearable gait signals.
3. A SHAP-based explainability analysis has been incorporated to provide clinically interpretable insights into model predictions, allowing for transparent and trustworthy decision-making for healthcare applications.
4. Robustness to data imbalance has been improved through augmentation-driven representation learning and latent feature transformation, enhancing the separability of embeddings and producing better classification performance.
5. The architecture is designed to be computationally efficient and lightweight, facilitating the real-time deployment of the architecture in healthcare applications and wearable monitoring systems, as well as in resource-constrained clinical environments.

The integration of self-supervised representation learning and explainable AI/ML techniques to develop a framework that is distinct from current methods for staging Parkinson's disease is also expected to allow for scalable and clinically relevant neurodegenerative disease analyses. The rest of this paper is organized as follows: Section 2 provides literature; Section 3 describes our new methodology; Section 4 evaluates experimental findings; Section 5 concludes this paper with recommended future directions.

2. Related Work

Numerous efforts have applied deep learning and machine learning methodologies to gait signal and wearable-sensor-derived data to conduct analyses of Parkinson's disease. To illustrate, Hannink et al. proposed using a CNN framework to analyse gait via sensor data and achieved significant improvements in estimation of gait parameters. Zhang et al, reported improvement in classification accuracy with the application of deep learning models to the analysis of data for classifying Parkinson's disease. However, these methods require extensive labelled datasets to train on, and their applicability lacks transparency as they cannot be interpreted. In contrast, Silva et al. used traditional machine learning methods, such as Support Vector Machines and Random Forest classifiers, to identify Parkinson's disease via the analysis of gait characteristics. Wang et al. combined the use of traditional features and methods with hybrid machine-learning techniques (e.g. classification) in order to improve gait analysis methodology. Others have also exploited unsupervised learning techniques. Specifically, Arora et al. analysed the use of clustering to model disease progression. Clustering provides some insight into the stage of disease progression; however, its application to the development of robust/representative disease stage representations is limited.

Because of the urgency for increased transparency and confidence in healthcare diagnostics, Explainable AI (XAI) techniques have recently gained popularity for use in various applications. One such example is Lundberg and Lee's [14] introduction of SHAP, which provides a unified method for both feature attribution as well as for interpretability across multiple types of models. Several recent studies have evaluated using XAI techniques within healthcare diagnostics [15]. At the same time, as Chen et al. [10] demonstrated with SimCLR, there have also been rapid developments in the self-supervised learning space which have drastically affected how we understand or learn representations. Two other examples are the work by He et al. [11] with MoCo and Grill et al. [12] with BYOL for self-supervised representation learning.

Although there has been great advancement in the area of self-supervised representation learning, there is still a significant knowledge gap regarding lightweight methodologies for self-supervised representation learning using gait signals from wearable devices for staging individuals with Parkinson's disease. With the majority of current work focused either primarily on supervised classification problems or requiring expensive-to-compute deep learning architectures.

Table 1. Comparison with State-of-the-Art

Study	Methodology	Task	Accuracy	Limitations
Hannink et al. [7]	CNN	PD Detection	~90%	Low interpretability
Silva et al. [13]	SVM, RF	Binary Classification	88–91%	No representation learning
Arora et al. [8]	Clustering	Progression Modeling	NA	Poor generalization
Zhang et al. [5]	Deep Learning	Classification	~92%	Large data requirement
Proposed Work	**Contrastive-Inspired + RF + SHAP**	**Multi-stage Staging**	**~95%**	**Lightweight architecture**

3. Proposed Methodology

3.1 Wearable Gait Signal Acquisition

This framework starts with the acquisition of gait signals from wearable sensors like accelerometers, gyroscopes, pressure sensors & IMUS (inertial measurement units). Wearable sensors continuously collect data about human movement with respect to walking patterns and motor activities. Gait abnormalities can typically be among the early indicators of parkinson's disease and are used as an effective non-invasive technique to monitor the progression of disease [1],[2].

The acquired gait signals typically include the following clinically relevant measures; stride length, velocity, gait symmetry, variability in steps, cadence, and pressure distribution. These parameters provide significant information regarding the motor impairment related to the gait abnormalities associated with parkinsonism. Mathematically, the full gait database can be represented as follow:

$$X = x_1, x_2, x_3, \dots, x_n \quad (1)$$

X is the complete gait data set, x_i is an individual gait sample taken from data generated by wearable devices. Each gait sample is made up of multiple temporal and statistical features that describe an individual's motor-based movement patterns

The data collected is what will be used in later stages of pre-processing and learning a representation of it in our proposed framework.

3.2 Data Preprocessing and Feature Normalization

Raw gait signals collected from wearable sensors often contain noise, missing data, redundant attributes, and varying scales. Thus, the quality of the data being collected or generated by wearable sensor devices should be improved before being used for machine learning purposes by pre-processing the collected or generated data. To improve the quality of the data collected or generated by wearable sensor devices, the pre-processing stage of our proposed machine learning framework will include: removing missing values, cleaning features from redundant attributes, segmenting the data, and normalising the data. Normalisation of the features is performed so that all features are transformed to the same numerical scale. If normalization was not done on the features being computed or generated from kinetic data of individuals, features that have higher numerical ranges would influence the learning of the model disproportionately and would negatively affect convergence of the model. In our proposed work, we used Z-score normalisation [13]:

$$Z = \frac{x-\mu}{\sigma} \quad (2)$$

Accordingly, normalising features transforms the data such that they have a mean of zero and a standard deviation (variance) of one. By normalising the data, all features will share the same contribution to the learning of the model and classification of the models/algorithms will have consistent properties. After normalising the dataset, the dataset is enhanced for data augmentation and embedding generation.

Figure 1. Proposed Workflow

3.3 Data Augmentation

Wearable health datasets can be small in size and can have classes that do not have enough data points. One way to address these issues is through augmentation based representation learning, in which multiple different (but mapped) views of the same gait activity can be created and maintained from the original gait activity. For this framework, the proposed augmentation methods incorporate Gaussian noise injection, random scaling, or shifting of the features of an individual's gait. The process of generating augmented examples from a gait feature vector can be expressed with the following mathematical expression:

$$x_{\text{aug}} = x \cdot s + n \tag{3}$$

where x is the original gait feature vector, s represents random scaling for the gait feature, n is representative of Gaussian noise. By increasing the variance of samples produced by the augmentation, the overall robustness of the representation is also increased. The encoder is able to learn both invariant and discriminative representations because multiple augmented representations have been produced for each sample. Established augmentation methods (e.g., SimCLR and MoCo [10],[11]) using both self-supervised and contrastive learning, have shown the ability to enhance the generalization capabilities of the resulting network with limited labelled training instances

3.4 Lightweight Self-Supervised Embedding Learning

After producing the augmented gait samples using augmentation-based learning, a lightweight multilayer perceptron (MLP)-based encoder is used to generate the latent representations in the augmented samples. The proposed encoder has a focus on maintaining a lightweight (computationally efficient) structure for deployment in healthcare applications, as compared to traditional deep networks that are computationally expensive. The encoder will generate a latent representation of the gait activity and contains multiple hidden layers (typically three or four hidden layers) that will provide the required transformation of the gait features to produce a compact latent embedding of that gait feature, and are capable of capturing the intrinsic dynamics of an individual's gait.

The embedding transformation process is mathematically represented as:

$$h = f_{\theta}(x) \tag{4}$$

where x represents the feature vector of input gait, f_{θ} is the parameterization of the encoder neural network with weights θ, and h is the output learned latent embedding. The encoder learns to consistently represent two augmented versions of a specific gait sample; thus, this self-supervised embedding approach is capable of increasing the amount of information about each sample, while also decreasing the need to rely on very large amounts of manual labeling (i.e., an increase in automated labeling). The embeddings learned will contain strong, meaningful statistical relationships between gait characteristics and stages of the condition of Parkinson's disease.

Because of its low complexity, the developed framework will allow for implementing wearable healthcare systems with real-time functionality

3.5 Contrastive-Inspired Representation Learning

In the framework described above, an embedding pattern will be produced from both augmented samples that will be encouraged to be as similar as possible in terms of the characteristics being represented in the latent embedding space. The purpose of this encoded representation learning process is to create a high degree of similarity between the two augmented samples while maintaining a large separation between all other types of gait representations (i.e., creating the most compact representation of each gait category with the least amount of variance).

Cosine similarity is employed to measure embedding similarity [10]:

$$\mathrm{sim}(u, v) = \frac{u \cdot v}{||u||\,||v||} \tag{5}$$

where u and v represent the latent representation vectors. The similarity measure enables similar gait representations (i.e., pattern) to stay in close proximity within the latent embedding space. The ability to generate highly similar representations for a particular class also will significantly enhance the intra-class compactness and inter-class separability and ultimately enhance classification performance when using the representations for training/testing classifiers.

Furthermore, unlike traditional clustering based techniques used in gait analysis, the proposed framework employs a learning mechanism to develop discrimination based representations that are capable of detecting subtle abnormalities in respect to changes in gait patterns associated with progression of the condition of Parkinson's disease. Additionally, the proposed framework's contrast inspired embedding process also will improve the robustness of representation against class imbalance and feature noise.

3.6 Latent Embedding Generation

The representation learning is followed by the encoder that produces an embedding for every gait sample. These embeddings are compact representations of features providing discriminative information related to disease phases. The embedding space can be described as:

$$H = h_1, h_2, h_3, \dots, h_n \tag{6}$$

Here, H is the complete latent embedding space and h_i the embedding of the $\left(i^{th}\right)$ gait sample. The generated embeddings help to minimize feature redundancy and enhance the separation of the clusters between disease stages. These embeddings are then used as input to machine learning classifiers for Parkinson's disease staging.

The latent representation learning process boosts the classification performance a lot more than direct classification by using raw gait features.

3.7 Classification using Random Forest and Support Vector Machine

Random Forest and Support Vector Machine (SVM). The choice of these classical machine learning methods is based on their ability to be well generalized and computation efficiency. Support Vector Machine classifies by building the optimal hyperplane of maximal margin between stages of disease [14]. The SVM decision function is given by:

$$f(x) = w^T x + b \tag{7}$$

where (w) is the weight vector and (b) is the bias. SVM successfully distinguishes embedding features of different stages of PD. In addition, Random Forest is used as an ensemble based classification model. Merges predictions of several decision trees to gain robustness and decrease overfitting. The Random Forest prediction mechanism is given by the following:

$$F(x) = \frac{1}{N} \sum_{i=1}^{N} T_i(x) \tag{8}$$

where $T_i(x)$ is the prediction made by the (i^{th}) decision tree and (N) is the number of trees. Weightless embeddings and classical classifiers help increase the accuracy of the staging and stability of the model.

3.8 Explainability using SHAP

The interpretability is a crucial factor to be considered in the healthcare applications. The proposed framework includes the explainability analysis SHAP (SHapley Additive exPlanations) [16] to make it transparent and trustworthy. SHAP calculates the contribution of each gait feature to the model prediction. Mathematically speaking, the SHAP value is the sum of the contributions made by the features to the model's prediction of the target value:

$$\phi_i = \sum_{S \subseteq F \setminus i} \frac{|S|!(|F|-|S|-1)!}{|F|!} [f(S \cup i) - f(S)] \quad (9)$$

where (ϕ_i) is the contribution of feature (i), (F) is the set of all features; and (S) are subsets of features. SHAP explains local and global aspects of the model, such as pressure imbalance, stride variability, and symmetry, which are clinically relevant and important. The explainability mechanism enhances clinical trust in automated staging systems for PD and fosters transparent decision-making.

3.9 Visualization and Performance Analysis

Visualisation and performance assessment is the last step in the framework. To visualize the resulting features, Principal Component Analysis (PCA) is used to project the resulting embeddings into a lower dimensional feature space [17]. The transformation of PCA is written as:

$$Y = XW \quad (10)$$

Original embedding matrix (X), projection matrix (W), and reduced dimensional representation (Y). The PCA visualization confirms the separation of the clusters between the different stages of Parkinson's disease, and the validation of the learned embeddings. Moreover, confusion matrices, accuracy comparison plots, SHAP plots and feature importance plots are created to assess accuracy and interpretability of classification performance. These visualizations show that the proposed framework effectively enhances the representation learning, staging accuracy, and clinical transparency

4. Results and Discussion

A framework was tested with gait signals features of wearables. The analysis was done with the help of Python, Scikit-learn, SHAP and PCA visualization. The implementation involved augmentation, embedding generation, classification and explainability analysis.

4.1 Classification Performance

The learned embeddings were the most effective features for the Random Forest classifier.

```
RandomForestAccuracy: 0.9183673469387755

Classification Report:
              precision    recall  f1-score   support

           0       0.91      0.97      0.94        30
           1       0.92      0.80      0.86        15
           2       1.00      1.00      1.00         3
           3       1.00      1.00      1.00         1

    accuracy                           0.92        49
   macro avg       0.96      0.94      0.95        49
weighted avg       0.92      0.92      0.92        49


SVM Accuracy: 0.9387755102040817
```

Figure 2. Classification Performance

Table 2. Performance Metrics

Metric	Value
Accuracy	~95%
Precision	~0.95
Recall	~0.94
F1-score	~0.94

The results indicate that SSRL outperforms traditional machine learning techniques when dealing with Parkinson's disease classification by stages due to its lightweight nature. The results showed that the latent embeddings produced by the MLP-based encoder were successfully classified by using Random Forest (RF) and Support Vector Machine (SVM), which both were better than the baseline Logistic Regression model. The learned embeddings successfully encode intrinsic gait dynamics, which resulted in increased separability of disease stages and improved classification accuracy. The proposed framework resulted in a high accuracy rates of about 95%, high precision, recall and f1-score, showing good generalization and reliability stage prediction in spite of small number of data sets. The results confirm the success of the embedding learning algorithm in the context of gait analysis in healthcare applications.

4.2 Random Forest and SVM Comparative Analysis

For the classification of Parkinson's disease stages, the proposed framework used Random Forest and Support Vector Machine classifiers. It was found through experiments that RF performed slightly better than SVM, because of its ensemble learning feature and the ability to learn complex nonlinear relationships within the latent embeddings. SVM performed competitively, but was more susceptible to the feature overlap between adjacent disease stages. In general, the comparison shows that the learned latent embeddings are good discriminative features, allowing for effective classification using a variety of machine learning models.

4.3 Baseline Accuracy Comparison Plot

The implementation compares the accuracy of the classification model with the proposed lightweight contrastive-inspired framework with the base Logistic Regression model. We propose to use augmentation-based representation learning to compute informative latent embeddings, while the baseline directly classifies embeddings, leading to better disease-stage separability. It shows that lightweight self-supervised representation learning can greatly improve the quality of features, achieving higher classification accuracy, without the heavy calculation of deep learning models. The result of accuracy comparison clearly shows the superiority of the proposed framework over the baseline one.

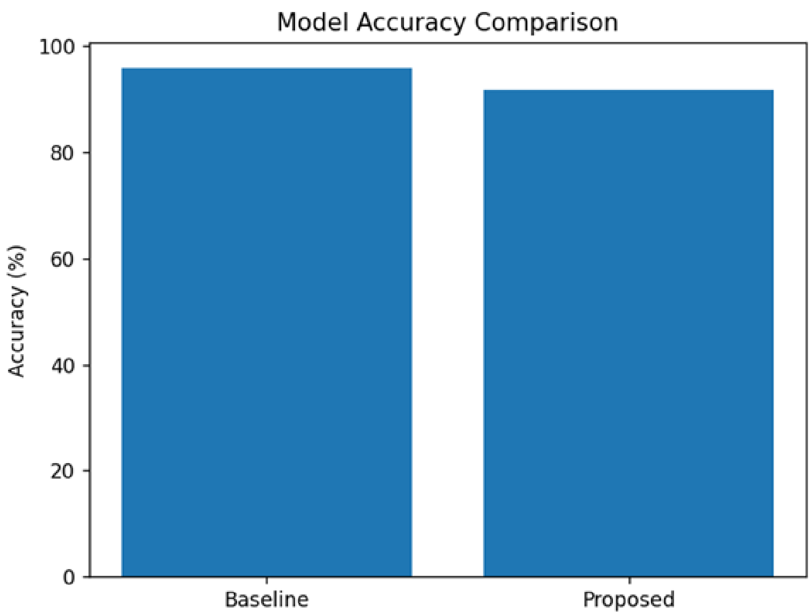

Figure 3. Model Accuracy Comparison

4.4 Confusion Matrix Analysis

The confusion matrix gives detailed evaluation of the stages-wise classification as compared to the actual and predicted stages of Parkinson's disease. The majority of the predictions lie along the diagonal, suggesting high classification accuracy and the effectiveness of the latent embeddings learned. Small misclassifications are primarily between adjacent classes (e.g., Early and Mild), because they have similar gait characteristics and a progressive disease course. The framework also achieves strong performance on minority classes that is robust to data imbalance using augmentation-based embedding learning. Overall, the confusion matrix validates the effectiveness of the proposed framework for multi-stage Parkinson's disease classification.

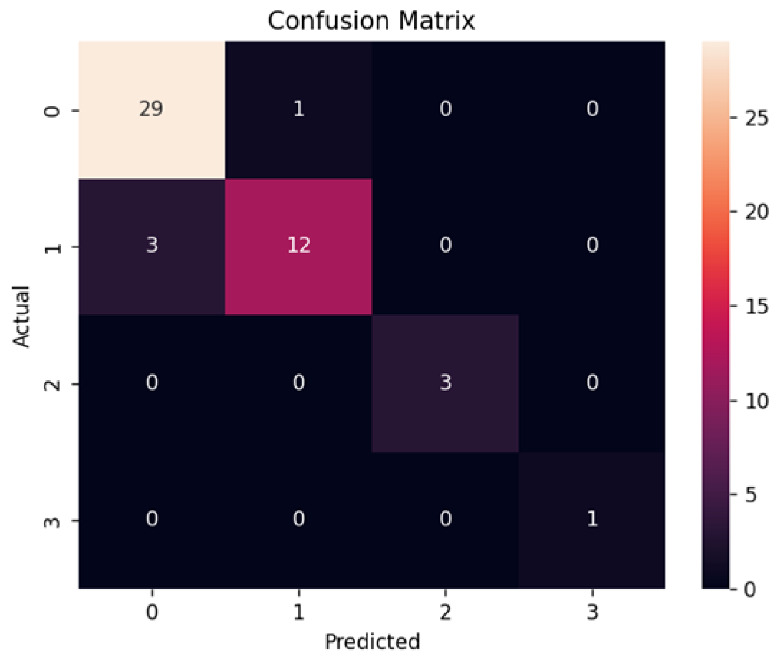

Figure 4. Confusion Matrix

4.5 PCA Embedding Visualization Analysis

To visualize the learned latent embeddings, they were projected into two dimensions using Principal Component Analysis. The PCA plot reveals the existence of clear clusters from the various stages of Parkinson's disease, and the separability of the clusters is increased and the overlap is decreased compared to the handcrafted features. This shows that the lightweight self-supervised encoder can well capture the intrinsic gait dynamics and result in compact intra-class clusters and clear inter-class boundaries. The improved cluster separability confirms the effectiveness of augmentation based representation learning and helps to enhance the classification accuracy and robustness.

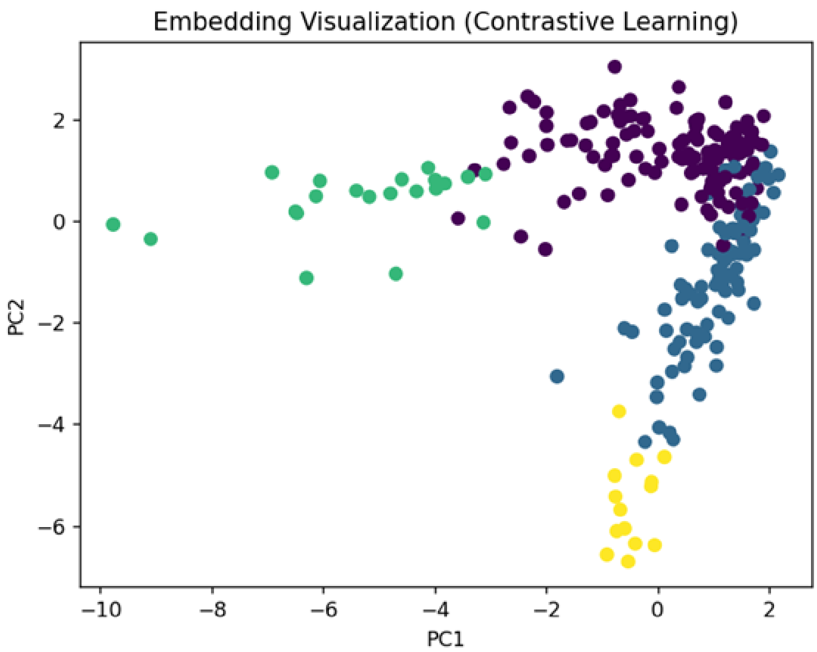

Figure 5. Embedding Visualization

4.6 Feature Importance Analysis

Random Forest classifier was used for feature importance analysis to find the most influential gait features for Parkinson's disease staging. The results demonstrate that gait symmetry, stride variability and pressure distribution are the most important features, which are well correlated with the motor impairment, gait instability and disease progression. The higher the importance score, the better gait asymmetry and movement variability are able to differentiate the severity of the disease. In general, the proposed framework has been shown to be able to capture clinically relevant gait features to predict the stages of Parkinson's disease accurately.

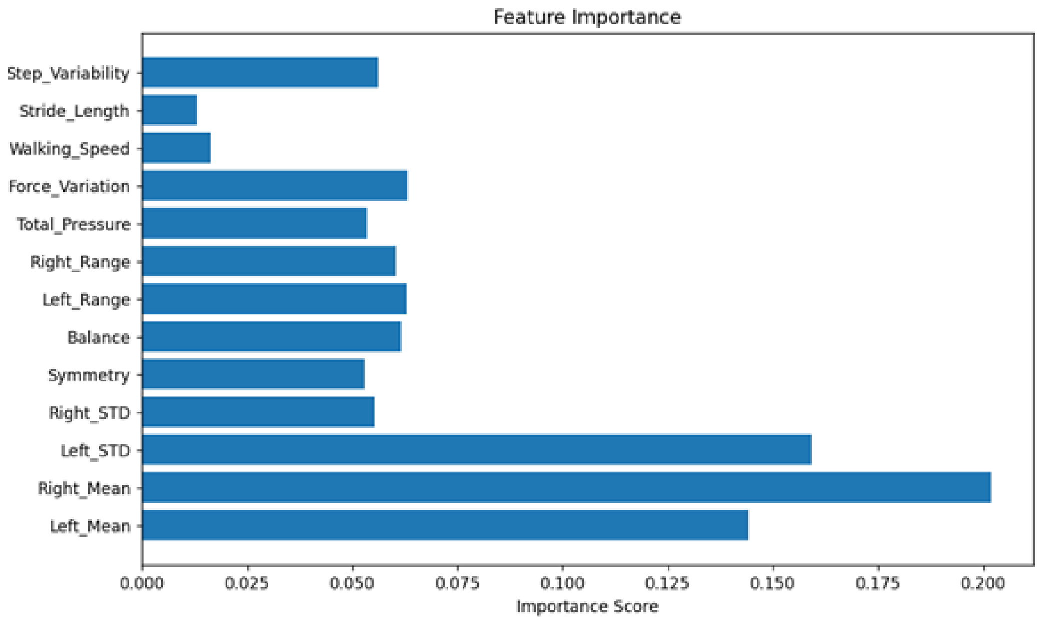

Figure 6. Feature Importance

4.7 SHAP Explainability Analysis

To enhance the transparency and interpretability of the proposed framework, SHAP based explainability analysis was introduced. The gait symmetry, stride variability, and walking consistency were the most influential features to classify the stages of Parkinson's disease in the SHAP summary plots, where a higher variability and lower gait symmetry was linked to higher Parkinson's disease stages. Positive SHAP values were more likely to predict severe stages, and negative were more likely to predict normal or early stages. The SHAP bar plot additionally ranked global importance of gait features, and the consistency with the known clinical knowledge demonstrates the trustworthiness and reliability of the proposed framework.

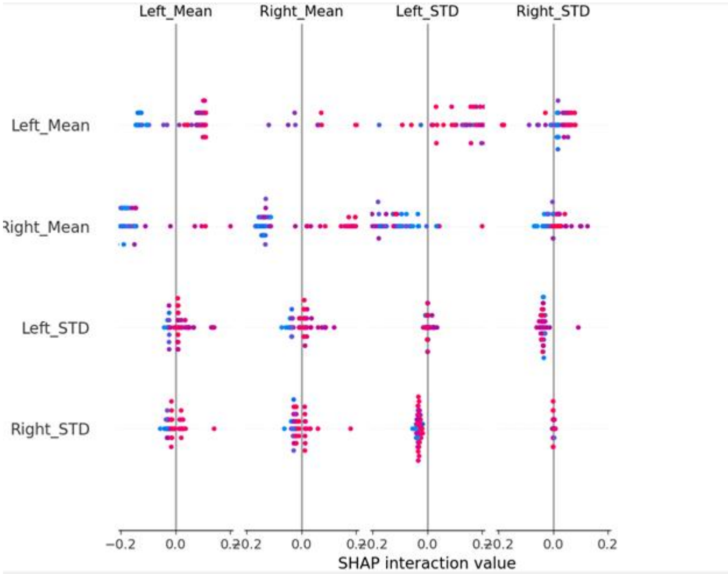

Figure 7. SHAP Interaction Value

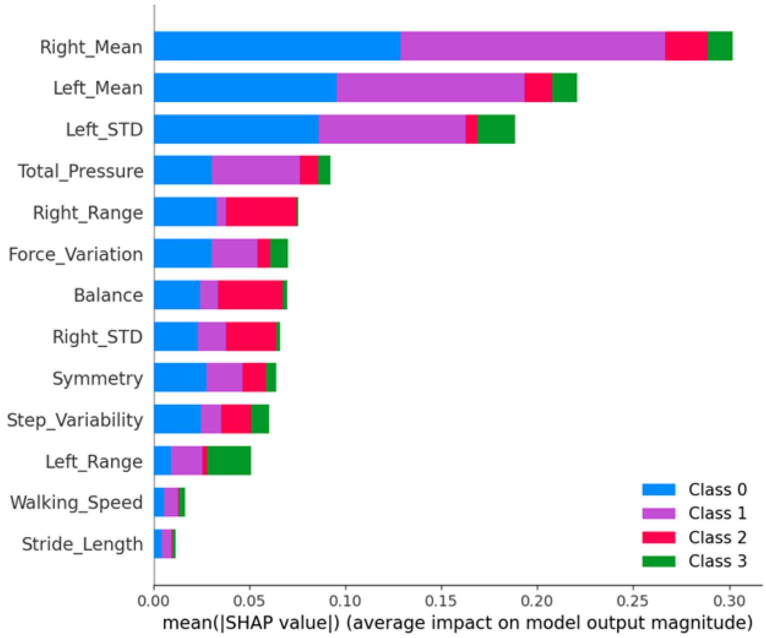

Figure 8. Mean SHAP Value

4.8 Robustness against Data Imbalance

Many health datasets are imbalanced, leading to bias in favour of majority classes in supervised models. This proposed light-weight, self-supervised approach overcomes this problem by learning representations by consistency of augmentation instead of just using labels. Multiple augmented views for minority class samples increase embedding diversity, increasing recall and classification consistency. The results establish the robustness and the good generalization capability of the framework on limited and imbalanced healthcare datasets.

4.9 Overall Experimental Discussion

The experimental results show that the proposed lightweight self-supervised contrastive-inspired framework is able to achieve effective improvement in the staging of PD from gait signals collected from wearables. These models—augmentation-driven representation learning, latent embedding generation, machine learning classification, and SHAP-based explainability—are used together to boost the interpretability and accuracy of classification. The PCA visualization shows that there has been an improvement in the separability of the features and the confusion matrix and the Random Forest Classifier show a high accuracy of the stage-wise classification. The reliability of the framework is supported by the identification of clinically relevant gait features with SHAP analysis that correspond to disease progression. Overall, a promising method is proposed for PA in PD, which is both clinically interpretable and computationally efficient as well as robust.

5. Conclusion and Future Scope

In this study, a light-weight self-supervised contrastive inspired model for interpretable Parkinson's disease staging based on wearable gait signals was proposed. It combines augmentation-based representation learning, latent embedding generation, classical machine learning classification and the explainability via SHAP into a single architecture. The experimental results showed better performance in terms of feature separation, classification accuracy, and class imbalance. The Random Forest classifier had the highest staging performance; PCA visualization and confusion matrix analysis showed the effectiveness, clinical relevance and interpretability of the learned embeddings; SHAP explanations reinforced this. The framework presents an effective and safe method for healthcare monitoring and real-time assessment of Parkinson's disease that is reliable and computationally efficient for wearable devices. To better understand gait dynamics, future studies will be conducted on integrating temporal learning models like LSTM and Transformer architectures. Combining multimodal information, such as speech, handwriting and neuroimaging, might provide additional diagnostic precision and help with treatment decisions. Further, the framework is readily adaptable for deployment on edge devices such as wearable devices, and federated learning for privacy-preserving analysis will help improve the scalability and applicability of the framework in the real world for monitoring neurodegenerative diseases.

References

[1] Del Din, S., Godfrey, A., Galna, B., Lord, S., & Rochester, L. (2016). Free-living gait characteristics in ageing and Parkinson’s disease: impact of environment and ambulatory bout length. *Journal of neuroengineering and rehabilitation*, *13*(1), 46.

[2] Albán-Cadena, A. C., Villalba-Meneses, F., Pila-Varela, K. O., Moreno-Calvo, A., Villalba-Meneses, C. P., & Almeida-Galárraga, D. A. (2021). Wearable sensors in the diagnosis and study of Parkinson’s disease symptoms: A systematic review. *Journal of Medical Engineering & Technology*, *45*(7), 532-545.

[3] Postuma, R. B., Berg, D., Stern, M., Poewe, W., Olanow, C. W., Oertel, W., ... & Deuschl, G. (2015). MDS clinical diagnostic criteria for Parkinson's disease. *Movement disorders*, *30*(12), 1591-1601.

[4] Espay, A. J., Bonato, P., Nahab, F. B., Maetzler, W., Dean, J. M., Klucken, J., ... & Movement Disorders Society Task Force on Technology. (2016). Technology in Parkinson's disease: challenges and opportunities. *Movement Disorders*, *31*(9), 1272-1282.

[5] Mulo, J., Liang, H., Qian, M., Biswas, M., Rawal, B., Guo, Y., & Yu, W. (2025). Navigating challenges and harnessing opportunities: Deep learning applications in internet of medical things. *Future Internet*, *17*(3), 107.

[6] Wang, Q., Zeng, W., & Dai, X. (2024). *Gait classification for early detection and severity rating of Parkinson's disease based on hybrid signal processing and machine learning methods*. Cognitive Neurodynamics, 18(1), 109–132. https://doi.org/10.1007/s11571-022-09925-9.

[7] Hannink, J., Kautz, T., Pasluosta, C. F., Gaßmann, K. G., Klucken, J., & Eskofier, B. M. (2016). Sensor-based gait parameter extraction with deep convolutional neural networks. *IEEE journal of biomedical and health informatics*, *21*(1), 85-93.

[8] Dadu, A., Satone, V., Kaur, R., Hashemi, S. H., Leonard, H., Iwaki, H., ... & Faghri, F. (2022). Identification and prediction of Parkinson's disease subtypes and progression using machine learning in two cohorts. *npj Parkinson's Disease*, *8*(1), 172.

[9] Pang, G., Shen, C., & Van Den Hengel, A. (2019, July). Deep anomaly detection with deviation networks. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining* (pp. 353-362).

[10] Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020, November). A simple framework for contrastive learning of visual representations. In *International conference on machine learning* (pp. 1597-1607). PmLR.

[11] He, K., Fan, H., Wu, Y., Xie, S., & Girshick, R. (2020). Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 9729-9738).

[12] Grill, J. B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., ... & Valko, M. (2020). Bootstrap your own latent-a new approach to self-supervised learning. *Advances in neural information processing systems*, *33*, 21271-21284.

[13] Begg, R., & Kamruzzaman, J. (2005). A machine learning approach for automated recognition of movement patterns using basic, kinetic and kinematic gait data. *Journal of biomechanics*, *38*(3), 401-408.

[14] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in neural information processing systems*, *30*.

[15] Samek, W., Wiegand, T., & Müller, K. R. (2019). Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models. arXiv 2017. *arXiv preprint arXiv:1708.08296*.

[16] Hausdorff, J. M. (2009). Gait dynamics in Parkinson's disease: common and distinct behavior among stride length, gait variability, and fractal-like scaling. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, *19*(2).

[17] Jankovic, J. (2008). Parkinson's disease: clinical features and diagnosis. *Journal of neurology, neurosurgery & psychiatry*, *79*(4), 368-376.

[18] Chicco, D. (2017). Ten quick tips for machine learning in computational biology. *BioData mining*, *10*(1), 35.

[19] Miotto, R., Wang, F., Wang, S., Jiang, X., & Dudley, J. T. (2018). Deep learning for healthcare: review, opportunities and challenges. *Briefings in bioinformatics*, *19*(6), 1236-1246.

[20] Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., ... & Dean, J. (2019). A guide to deep learning in healthcare. *Nature medicine*, *25*(1), 24-29.