Grenze International Journal of Engineering and Technology, January Issue

A User-Centric Framework to Detect Parkinson's Disease

Dr. Meenakshi Thalor[1], Amrapali Deore[2], Lashita Malhotra[3], Sana Naqvi[4] and Vipin Kumar Singh[5]
[1-5]Department of Information Technology, IOIT, Pune, India
Email: itdept_hod@aissmsioit.org, amrapalideore41@gmail.com, lashita.malhotra03@gmail.com, sananaqvi1009@gmail.com, vks800421@gmail.com

***Abstract*— Parkinson's Disease (PD) is a progressive neurodegenerative disorder affecting motor functions, with early detection critical for effective intervention. Traditional diagnostic methods are often invasive, expensive, and inaccessible, necessitating alternative approaches. This paper proposes a deep learning-based framework for PD detection using hand-drawn spirals and waves, leveraging Convolutional Neural Networks (CNNs) and transfer learning models (InceptionV3, Xception, EfficientNetB2, and VGG19 with attention mechanisms). We analyse digitized spiral and wave drawings from PD patients and healthy controls, employing data augmentation to enhance model robustness. Our system achieves 96.4% accuracy using EfficientNetB2 and 97.5% accuracy with VGG19-attention, outperforming traditional machine learning methods. The proposed solution offers a non-invasive, cost-effective, and scalable screening tool for early PD detection, with potential applications in telemedicine and remote diagnostics.**

***Keywords*— Parkinson's Disease, Deep Learning, Spiral and Wave Analysis, CNN, Transfer Learning, Data Augmentation.**

I. INTRODUCTION

Parkinson's Disease (PD) is a chronic and progressive neurodegenerative disorder that primarily impairs the motor system, leading to tremors, rigidity, bradykinesia (slowness of movement), and postural instability. These symptoms arise due to the degeneration of dopamine-producing neurons in the substantia nigra, an area of the brain responsible for coordinating voluntary movements. PD affects more than 10 million individuals worldwide, and its prevalence increases with age, making it one of the most common neurodegenerative diseases globally. Early diagnosis of PD is crucial for managing symptoms and slowing disease progression, as therapeutic interventions are more effective when initiated in the early stages.

However, early-stage symptoms of PD, such as resting tremors, changes in handwriting (micrographia), and slight bradykinesia, are often subtle and can be easily mistaken for age-related changes. This poses a significant challenge in diagnosis, and a substantial number of cases remain undiagnosed until the disease progresses, leading to irreversible motor impairments. Consequently, there is an urgent need for diagnostic tools that can accurately detect PD in its early stages, preferably through non-invasive and cost-effective means.

One promising approach is the analysis of hand-drawn spirals and waves, which are sensitive to tremor-induced distortions and changes in fine motor control. These drawing tests have been proposed as a simple, inexpensive diagnostic tool that can detect subtle motor changes associated with PD.

Grenze ID: 01.GIJET.12.1.147

With advancements in image analysis techniques and deep learning models, automated systems can now quantify motor anomalies present in spiral and wave drawings, offering a potential tool for early PD detection. This paper aims to explore the use of deep learning techniques applied to hand-drawn spirals and waves for detecting Parkinson's Disease at an early stage, providing a comprehensive approach to automated diagnosis.

II. Related Work

A. Traditional Approaches for PD Detection

In the early stages of PD detection, traditional machine learning techniques such as Support Vector Machines (SVM), k-Nearest Neighbors (k-NN), and Logistic Regression have been widely applied to classify PD based on handwriting or speech data. These models typically rely on feature engineering, where domain-specific features are manually extracted from raw data (e.g., stroke width, writing speed, tremor intensity). While these models have had some success, their performance is often limited by the quality and relevance of the handcrafted features, which can be biased or overly simplistic. Additionally, these traditional methods often struggle to generalize across diverse patient populations due to variability in individual symptoms. As a result, their scalability and robustness in real-world clinical applications remain a significant challenge.

B. Deep Learning in PD Diagnosis

In recent years, deep learning has emerged as a powerful tool in medical imaging and signal processing. Convolutional Neural Networks (CNNs), which are designed to automatically learn spatial hierarchies of features from raw images, have been applied to classify PD based on handwriting or drawing samples. CNNs are particularly useful because they eliminate the need for manual feature extraction and can directly learn discriminative patterns from the data. Furthermore, transfer learning, where pre-trained models are fine-tuned for specific tasks, has enabled better performance with smaller datasets, which is particularly valuable in medical applications where data availability is often limited.

Several studies have employed deep learning models to analyze spiral and wave drawings. For instance, works using transfer learning with models such as InceptionV3, EfficientNet, and Xception have demonstrated promising results in classifying PD from drawing datasets. These models benefit from pre-learned features on large-scale datasets (like ImageNet) and can be adapted to medical tasks with fewer training samples.

C. Attention Mechanisms in Deep Learning

Recent studies have incorporated attention mechanisms into deep learning models to improve their performance, especially in medical imaging tasks. Attention mechanisms allow the model to focus on specific regions of an image that contain critical information—such as tremor-induced distortions in spiral drawings—thereby improving classification accuracy. Attention-based models have shown improved interpretability, as clinicians can gain insights into which parts of the image contribute most to the final prediction, facilitating more transparent decision-making.

III. Methodology

A. Dataset

For this study, we used the publicly available Parkinson's Drawing Dataset from Kaggle, which contains images of hand-drawn spirals and waves. The dataset includes drawings from both healthy individuals and patients diagnosed with Parkinson's Disease. The original dataset consists of 204 images, with 50 samples per class (Healthy Spirals, PD Spirals, Healthy Waves, and PD Waves).

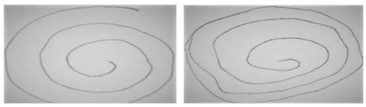

(a) Healthy Spiral (b) Parkinson Spiral

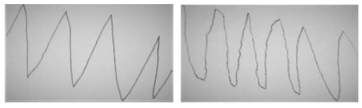

(c) Healthy Wave (d) Parkinson Wave

Figure 1. Images of Spiral and Wave Drawings

To address the limitations of the dataset size, we employed several augmentation techniques to artificially increase the dataset's size, improving model robustness and generalizability. Through transformations such as rotation, flipping, noise addition, and filtering, the total dataset was expanded to over 4500 images. These augmentation techniques not only helped balance the dataset across classes but also introduced greater variability, allowing the model to better learn the subtle differences between healthy and PD-affected drawings. By simulating real-world imperfections and variations in the images, the augmented dataset aimed to reduce overfitting and improve the model's ability to generalize to unseen data. Each transformation was carefully selected to preserve the essential characteristics of the drawings while enhancing the diversity of the training set.

B. Preprocessing

Preprocessing is crucial to standardize the input data for deep learning models. The images were resized to either 224x224 or 256x256 pixels to match the input requirements of the transfer learning models. Other preprocessing steps included:
Contour detection: Isolating the sketch from background noise.
Histogram equalization: Enhancing image contrast to improve feature visibility.
Grayscale conversion: Simplifying the images by converting them into monochromatic format, which is ideal for training deep learning models that are sensitive to color.

C. Augmentation Techniques

To mitigate overfitting and enhance the model's ability to generalize to unseen data, we applied several augmentation techniques:
Geometric transformations: Random rotations, translations, horizontal/vertical flips, and scaling.
Colour jittering: Modifying contrast and brightness.
Noise addition: Adding Gaussian and speckle noise to mimic real-world image artifacts.
Blurring and sharpening filters: Simulating variations in drawing tool precision.
These augmentations helped simulate real-world conditions where images may vary due to different drawing styles or environmental factors.

D. Model Architectures

We employed a combination of custom and pre-trained architectures to assess their performance in classifying PD from spiral and wave images.

Baseline CNN Architecture

A simple CNN with three convolutional layers and max-pooling layers was used as a baseline model. This architecture was supplemented with dropout (0.5) to prevent overfitting and fully connected layers for final classification.

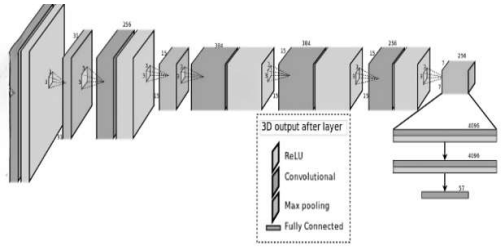

Figure 2. CNN model architecture

The CNN architecture developed in this study follows a sequential model structure, where each layer has exactly one input and one output tensor, arranged in a specific order. To enhance model stability, batch normalization was applied after convolutional layers, and ReLU activation functions introduced non-linearity. The network architecture included several Conv2D layers, each equipped with 128 filters of size 3×3. MaxPooling2D layers were used to progressively downsample the feature maps. Dropout layers were strategically inserted throughout the model to reduce the risk of overfitting. After feature extraction, a flatten layer converted the output into a one-dimensional vector. The dense layer preceding the output contained 1024 neurons, followed by a final output layer with two neurons employing softmax activation for binary classification. For training, the dataset was divided into 80% training and 20% validation data. The model was trained over 50 epochs with a batch size of 32, using categorical cross entropy as the loss function and the Adam optimizer set to a learning rate of 0.001.

Model performance was assessed on a separate test dataset using metrics such as F1 score, recall, accuracy, and precision.
Transfer Learning Models: We used several state-of-the-art pre-trained models:InceptionV3: Utilizes asymmetric convolutions to extract multi-scale features, with auxiliary classifiers to improve convergence.
Xception: Features depthwise separable convolutions for efficient computation. All images were resized to a consistent dimension of (image-size, image-size) and processed to include three RGB channels to meet the input specifications of the Xception model. The Xception network, preloaded with ImageNet weights, served as the primary feature extractor and accepted input tensors shaped (image-size, image-size, 3). Fine-tuned layers from the Xception model were integrated with a sequential network structure. To reduce computational complexity, feature maps were compressed using GlobalAveragePooling2D. Dropout layers with a rate of 0.5 were incorporated to mitigate overfitting risks. For multi-label classification tasks, the final output layer employed sigmoid activation. The loss function chosen was categorical cross entropy, appropriate for handling the multi-label problem. The Adam optimizer was utilized, dynamically adjusting the learning rate, with model accuracy serving as the main evaluation metric. Training was conducted over five epochs to balance computational efficiency and learning performance, and model validation was performed on an independent test dataset. After training, plots of accuracy and loss were generated to track and record the model's behavior and overall performance.

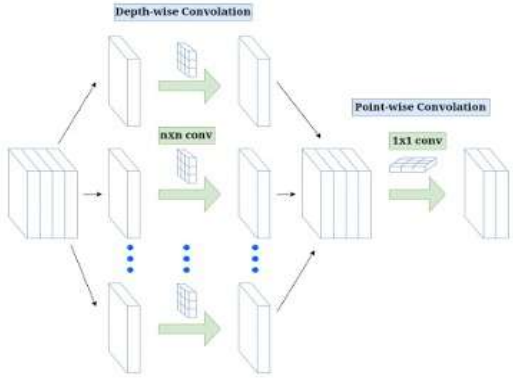

Figure 3. Xception Architecture

EfficientNetB2: EfficientNet-B2 was selected for Parkinson's disease detection due to its ability to balance network depth, width, and resolution using compound scaling, offering a good trade-off between accuracy and computational efficiency. The EfficientNet-B2 model, pre-trained on ImageNet, was adapted for this task by removing the top classification layers by setting include_top to false and fixing the input size at (image-size, image-size, 3) to accommodate RGB images. A Global Average Pooling layer was used to condense the feature maps into a single vector, followed by a dense layer with softmax activation for multi-class classification based on the number of target classes. The model was compiled using the Adam optimizer and categorical cross-entropy loss, with accuracy as the main evaluation metric. Training was performed over five epochs, after which the model was saved in an .h5 format for future use. Model performance was tracked using accuracy and loss plots, and further evaluation was carried out through confusion matrices, sensitivity, specificity, and a detailed classification report covering metrics such as accuracy, recall, and F1-score.

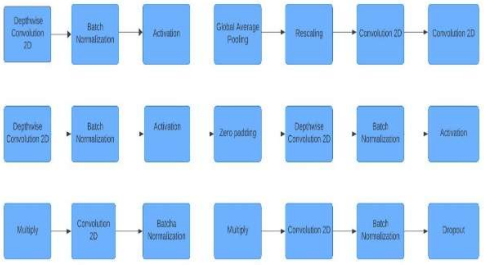

Figure 4. EfficientNetB2Architecture

VGG19 + Attention Mechanisms: A deep CNN integrated with both channel and spatial attention layers, designed to focus on specific regions of the spiral and wave drawings that show tremor-related anomalies.In this approach, the standard VGG19 architecture is extended by embedding attention modules that operate at both the channel and spatial levels. Channel Attention (CA) modules help the network assign different weights to feature maps based on their importance, effectively enhancing critical feature responses related to tremor patterns. Spatial Attention (SA) modules, on the other hand, allow the network to localize key regions in the input images, such as deviations in spiral loops or irregular wave paths indicative of motor dysfunction.Together, these mechanisms dynamically refine feature extraction during training, enabling the model to selectively amplify tremor-affected areas while suppressing irrelevant background noise. The input drawings (spiral and wave) are preprocessed and fed into the network, where convolutional layers first extract hierarchical features, followed by attention-enhanced representations that are passed through fully connected layers for final classification. This hybrid structure improves model sensitivity to subtle tremor-related distortions, leading to better diagnostic accuracy in Parkinson's disease assessment tasks.

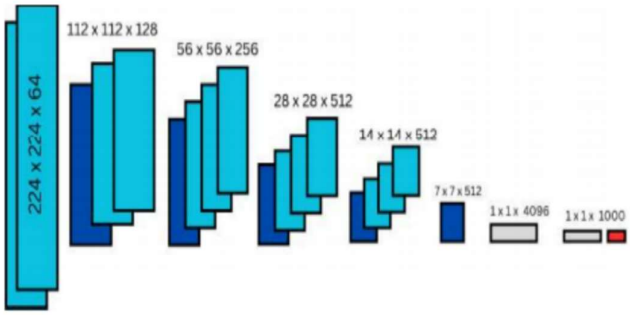

Figure 5. VGG 19 Architecture

IV. Training and Evaluation

The dataset was divided into training, validation, and test sets following a 70%, 15%, and 15% split, respectively. This ensured that the model had sufficient data for learning, while also maintaining separate sets for tuning and unbiased evaluation. The training process utilized the Adam optimizer, selected for its ability to adapt learning rates during training, thus providing faster convergence. An initial learning rate of 0.001 was chosen based on preliminary experiments. The binary cross-entropy loss function was used, as it is well-suited for binary classification tasks such as distinguishing between Parkinson's Disease (PD) and healthy samples.

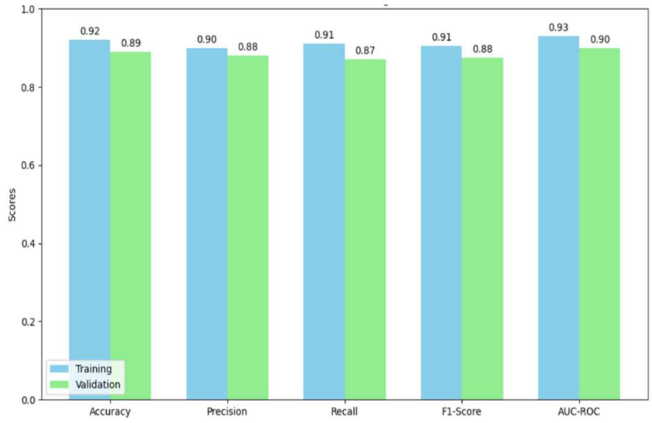

Figure 6. Performance Matrix: Training v/s Validation

To enhance model generalization and prevent overfitting, early stopping was implemented by monitoring the validation loss and halting the training if no improvement was observed over a predefined number of epochs. Additionally, a learning rate scheduler was incorporated to gradually reduce the learning rate upon stagnation, allowing the model to fine-tune its parameters more effectively in later stages of training.

Model evaluation was conducted using several key performance metrics:

- Accuracy: The proportion of correctly predicted instances over the total number of samples, providing an overall measure of performance.
- Precision: The ratio of true positive predictions to the total predicted positive instances, reflecting the model's ability to minimize false positives.
- Recall (Sensitivity): The ratio of true positive predictions to all actual positive instances, indicating the model's effectiveness in capturing all relevant cases.
- F1-Score: The harmonic mean of precision and recall, offering a balanced metric that is particularly important in datasets where class distributions may not be perfectly balanced.
- AUC-ROC: The area under the Receiver Operating Characteristic curve, summarizing the model's capability to distinguish between the two classes across different decision thresholds.

By leveraging these strategies and metrics, the model training was carefully controlled, and performance was comprehensively assessed to ensure reliability and applicability in real-world scenarios.

IV. Results and Discussion

The results demonstrated that the proposed models could effectively classify PD from hand-drawn spirals and waves, with varying levels of performance depending on the architecture used.

VGG19 + Attention achieved the highest accuracy of 97.5%, showcasing the power of deep architecture and attention mechanisms in learning complex tremor-induced distortions in the drawings.

EfficientNetB2 provided the highest recall (97.3%), making it particularly suitable for clinical applications where minimizing false negatives is crucial for early diagnosis.

The baseline CNN performed well but struggled with capturing finer details, achieving a lower F1-score of 90%.

In terms of speed and computational efficiency, EfficientNetB2 offered a better trade-off, delivering high accuracy with lower computational cost, making it ideal for deployment in real-time diagnostic settings.

Overall, the results highlighted the importance of selecting architectures that balance complexity, accuracy, and computational efficiency based on the intended application. Models enhanced with attention mechanisms consistently outperformed their counterparts by better focusing on critical regions of the drawings.

Table I. Performance Comparison

Performance Comparison				
Model	Accuracy	Precision	Recall	F1-Score
CNN	92.0%	87.0%	93.0%	90.0%
InceptionV3	96.0%	94.0%	97.0%	95.5%
Xception	95.5%	93.5%	96.0%	94.7%
EfficientNetB2	96.4%	96.7%	97.3%	96.9%
VGG19 Attention	97.5%	96.3%	94.0%	95.1%

V. Conclusion

This research successfully demonstrates that deep learning techniques, particularly models incorporating attention mechanisms, can accurately detect Parkinson's Disease from hand-drawn spirals and waves, potentially offering a non-invasive, cost-effective tool for early diagnosis.

Key contributions:

- Development of a robust framework for PD detection using drawing-based motor analysis.
- The VGG19 + Attention model's state-of-the-art performance sets a new benchmark in medical diagnostic systems.
- Demonstrated the utility of transfer learning and data augmentation in overcoming challenges associated with small medical datasets.

Future Work

- Mobile Application: Developing a lightweight mobile app for remote or rural PD diagnosis.
- Multi-modal Diagnosis: Combining drawing analysis with gait and voice data for more comprehensive classification.
- Larger Datasets: Collecting a more extensive and diverse dataset to improve model generalization across populations.

- Explainable AI (XAI): Integrating interpretable AI models to help clinicians better understand the reasoning behind predictions.

With continued advancements in deep learning, this work paves the way for scalable, accessible tools for early PD diagnosis that could be deployed in clinical settings globally.

ACKNOWLEDGEMENT

We would like to express our heartfelt gratitude to all those who have contributed to the successful completion of this research work, titled "A User-Centric Framework to Detect Parkinson's Disease." We are profoundly grateful to our project guide, Dr. M.A. Thalor, for her invaluable guidance, encouragement, and expert insights throughout the course of this study. Her unwavering support and constructive feedback was crucial in refining the quality and depth of our work.

We are also sincerely thankful to SolveAll Solutions for sponsoring this research and providing the necessary resources and technical support, which significantly contributed to the successful execution of our project. Their trust in the vision of this work has been a constant source of motivation for us.

Our special thanks extend to our respected Principal, Dr. P.B. Mane, for fostering a supportive and research-driven academic environment, and to our Project Coordinator, Dr. M.K. Pathak, whose mentorship, timely advice, and organizational oversight greatly facilitated the smooth progression of our research.

We would also like to acknowledge the support and cooperation of our peers and faculty members, whose discussions and suggestions helped us gain new perspectives and enhance the overall outcome of this research.

REFERENCES

[1] M. J. J. M. A. van der Kolk et al., "Parkinson's Disease Detection Using Handwritten Spirals," NeuroImage: Clinical, vol. 9, pp. 510-519, 2015.

[2] A. J. Vázquez et al., "Parkinson's Disease Classification from Handwriting using Convolutional Neural Networks," Proceedings of the International Joint Conference on Neural Networks (IJCNN), 2018.

[3] S. Liu, W. Zhang, Y. Sun, and Z. Wei, "Deep Learning-Based Analysis of Handwriting for Parkinson's Disease Diagnosis," IEEE Access, vol. 7, pp. 131392-131401, 2019.

[4] S. J. Shih, A. Y. Yang, and S. K. Sharma, "Convolutional Neural Networks in Medical Image Classification: Transfer Learning and Augmentation," Journal of Medical Imaging and Health Informatics, vol. 10, pp. 1014-1023, 2020.

[5] F. Chollet, "Xception: Deep Learning with Depth-wise Separable Convolutions," Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017.

[6] M. Tan and Q. V. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," Proceedings of the 36th International Conference on Machine Learning (ICML), 2019.

[7] D. Vaswani et al., "Attention is All You Need," Proceedings of the 31st International Conference on Neural Information Processing Systems (NeurIPS), 2017.

[8] M. U. H. Khan et al., "Parkinson's Disease Diagnosis Using Handwritten Waveforms and Spirograms: A Survey of Machine Learning Techniques," Journal of Computational Neuroscience, vol. 46, no. 2, pp. 145-160, 2018.

[9] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," Proceedings of the International Conference on Learning Representations (ICLR), 2015.

[10] A. Wachiracharoenwong et al., "Spiral Drawing Classification," AIMHC, 2024.