People's Democratic Republic of Algeria
Ministery of Higher Education and Scientific Research
Mohamed El Bachir El Ibrahimi University of Bordj Bou Arréridj
Faculty of Mathematics and Informatics
Informatics Department



DISSERTATION
Presented in fulfillment of the requirements of obtaining the degree
Master in Informatics
Specialty: Technologies de l'information et de la communication

THEME

An Explainable AI System for Parkinson's Disease Diagnosis from MRI Images

Presented by:

Abderraouf CHAA

Publicly defended on: 30/06/2026

In front of the jury composed of:

President:Pr. Abdelouahab ATTIA

Examiner: Dr Sara BOUTOUHAMI

Supervisor: Dr Mourad NOUIOUA

2025/2026

Dedication

To my beloved family,

whose unconditional love and endless support

have been my greatest strength throughout this journey.

To my friends,

for their encouragement and presence in every moment.

To all my professors,

who have guided and inspired me throughout my academic path.

This work is dedicated to all of you.

Acknowledgment

First and foremost, I would like to express my sincere gratitude to my supervisor, for their invaluable guidance, continuous support, and constructive feedback throughout the realization of this work. Their expertise and dedication have been essential to the successful completion of this thesis.

I would also like to extend my heartfelt thanks to all the professors of the Department of Computer Science at the University of Mohamed El Bachir El Ibrahimi of Bordj Bou Arréridj, whose teachings and knowledge have shaped my academic journey and provided me with the foundation necessary to carry out this research.

I am deeply grateful to my family for their unconditional love, patience, and moral support throughout my studies. Their encouragement has been a constant source of motivation.

I also wish to thank my friends for their support and for the moments of joy and solidarity we shared during this academic journey.

Finally, I would like to acknowledge the Kaggle platform for providing free access to GPU resources, which made the training of deep learning models possible within the scope of this work.

Abstract

Parkinson's disease, one of the most common neurodegenerative disorders worldwide, is still difficult to identify in its early stages with traditional clinical techniques. A deep learning-based method for automatically identifying Parkinson's disease from brain MRI pictures is presented in this thesis. The NTUA Parkinson dataset, which consists of over 43,000 brain MRI pictures categorized into PD and Non-PD classes, was used to refine the Xception convolutional neural network, which had previously been trained on ImageNet. The model was tested as a feature extractor in conjunction with two machine learning classifiers, Random Forest (81.40%) and XGBoost (79.13%), as well as as a stand-alone deep learning classifier with an accuracy of 84.19%.To ensure interpretability, three Explainable AI (XAI) techniques were applied: Grad-CAM for visual explanation, SHAP for quantifying feature importance, and a Grad-CAM-based segmentation approach to localize the brain regions most associated with the disease, particularly the Substantia Nigra and Basal Ganglia. Finally, a fully functional desktop application, ParkinsonAI, was developed using CustomTkinter to integrate detection, patient management, XAI visualization, and medical report generation, complemented by a commercial web platform built with PHP and JavaScript for distribution to researchers and healthcare institutions. Keywords: Parkinson's Disease, Magnetic Resonance Imaging (MRI), NTUA Dataset, Deep Learning, Convolutional Neural Networks (CNNs), Xception, Transfer Learning, Ran iv dom Forest, XGBoost, Explainable Artificial Intelligence (XAI), Grad-CAM, SHAP

Résumé

La maladie de Parkinson, l'un des troubles neurodégénératifs les plus répandus dans le monde, demeure difficile à diagnostiquer à un stade précoce avec les techniques cliniques traditionnelles. Ce mémoire présente un système basé sur l'apprentissage profond pour la détection automatique de la maladie de Parkinson à partir d'images d'IRM cérébrales. Le réseau de neurones convolutif Xception, pré-entraîné sur ImageNet, a été affiné sur la base de données NTUA Parkinson, qui comprend plus de 43 000 images d'IRM cérébrales réparties en deux classes : PD et Non-PD. Le modèle a été évalué à la fois comme classifieur d'apprentissage profond autonome, atteignant une précision de 84,19%, et comme extracteur de caractéristiques combiné à deux classifieurs de machine learning, Random Forest (81,40%) et XGBoost (79,13%). Afin de garantir l'interprétabilité, trois techniques d'Intelligence Artificielle Explicable (XAI) ont été appliquées : Grad-CAM pour l'explication visuelle, SHAP pour la quantification de l'importance des caractéristiques, et une approche de segmentation basée sur Grad-CAM afin de localiser les régions cérébrales les plus associées à la maladie, notamment la Substantia Nigra et les Ganglions de la Base. Enfin, une application de bureau entièrement fonctionnelle, ParkinsonAI, a été développée avec CustomTkinter afin d'intégrer la détection, la gestion des patients, la visualisation XAI et la génération de rapports médicaux, complétée par une plateforme web commerciale développée en PHP et JavaScript pour sa distribution auprès des chercheurs et des établissements de santé.

Mots-clés : Maladie de Parkinson, Imagerie par résonance magnétique (IRM), Jeu de don nées NTUA, Apprentissage profond, Réseaux de neurones convolutifs, Xception, Apprentis sage par transfert, Random Forest, XGBoost, Intelligence artificielle explicable, Grad-CAM, SHAP.

ملخص

يُعدّ مرض باركنسون أحد أكثر الأمراض التنكسية العصبية شيوعاًا في العالم، ولا يزال التشخيص المبكر يُمثّل تحدياًا سريرياًا هاماًا. تُقدّم هذه الأطروحة نظاماًا قائماًا على التعلّم العميق للكشف التلقائي عن مرض باركنسون من صور الرنين المغناطيسي الدماغي، وذلك في إطار شهادة الماجستير في تكنولوجيا المعلومات والاتصالات بجامعة محمد البشير الإبراهيمي ببرج بوعريريج.

يعتمد النظام المقترح على شبكة عصبية التفافية عميقة مدرّبة مسبقاًا وصُقلت على مجموعة بيانات تضمّ ما يقارب ثلاثاسوأربعين ألف صورة رنين مغناطيسي مُقسّمة إلى فئتين: مرضى باركنسون وغير المصابين بالمرض. في مرحلة لاحقة، استُخدمت الشبكة كمستخرج للميزات التي خُفِّضت باستخدام خوارزمية اختيار الميزات، ثم دُرِّب نموذجان للتعلم الآلي التقليدي على هذه الميزات.

ولضمان قابلية تفسير النتائج، طُبِّقت ثلاث تقنيات للذكاء الاصطناعي القابل للتفسير: التفسير المرئي لقرارات الشبكة، وتحديد أهمية الميزات، والتجزئة لتحديد مناطق الدماغ الأكثر ارتباطاًا بالمرض.

الكلمات المفتاحية: مرض باركنسون، التعلم العميق، الشبكات العصبية الالتفافية، التعلم بالنقل، الغابة العشوائية، تعزيز التدرج، الذكاء الاصطناعي القابل للتفسير، التصوير بالرنين المغناطيسي.

Table of contents

Arabic Abstract **v**

List of Figures **xi**

List of Tables **xiii**

General Introduction **xiv**

1 Parkinson's Disease and Diagnostic Methods **1**

1.1 Introduction 1

1.2 Definition and Impact of Parkinson's Disease 2

1.3 Diagnostic Methods of Parkinson's Disease 2

1.4 Magnetic resonance imaging (MRI) 3

1.4.1 Definition of MRI 3

1.4.2 Principles and Basic Concepts of MRI 4

1.4.3 MRI for the Diagnosis of Parkinson's Disease Symptoms 4

1.5 The Role of Artificial Intelligence in Diagnosis 6

1.6 Conclusion 6

2 Deep Learning Models and Machine Learning Algorithms for Parkinson's Detection **7**

2.1 Introduction 7

2.2 Machine Learning, Deep Learning and Supervised Classification 7

2.2.1 Machine Learning 7

2.2.2 Deep Learning 8

2.2.3 Advantages of Deep Learning over Machine Learning 8

2.2.4 Architecture of a Computer-Aided Diagnosis System for Parkinson's Disease . . . 9
2.3 Convolutional Neural Networks (CNNs) . . . 9
2.3.1 Feature Extraction Block . . . 10
2.3.2 Classification Block . . . 10
2.4 Transfer Learning and Pre-trained Models . . . 11
2.5 Transfer Learning Strategies . . . 11
2.5.1 Strategy 1: Full Training . . . 12
2.5.2 Strategy 2: Partial Training . . . 12
2.5.3 Strategy 3: Freezing all Convolutional Layers . . . 12
2.6 Transfer Learning Process . . . 12
2.6.1 Pre-trained Model Selection . . . 12
2.7 Xception . . . 12
2.7.1 Xception Architecture . . . 13
2.8 Classification Methods . . . 14
2.8.1 Random Forest (RF) . . . 14
2.8.2 XGBoost . . . 14
2.9 Explainable AI (XAI) . . . 15
2.9.1 Grad-CAM . . . 16
2.9.2 SHAP . . . 18
2.9.3 Segmentation . . . 21
2.10 Conclusion . . . 21

3 Proposed Methodology **23**
3.1 Introduction . . . 23
3.2 The Proposed System's General Architecture . . . 24
3.3 UML Diagrams . . . 25
3.3.1 Use Case Diagram . . . 25
3.3.2 Class Diagram . . . 26
3.3.3 Sequence Diagram . . . 27
3.4 Dataset Presentation . . . 27
3.4.1 Dataset Description . . . 27
3.4.2 Class Distribution . . . 28

3.5 Data Preprocessing . . . 28
3.5.1 Image Loading and Format Conversion . . . 28
3.5.2 Image Resizing . . . 28
3.5.3 Normalization . . . 28
3.5.4 Data Augmentation . . . 29
3.5.5 Dataset Splitting . . . 29
3.5.6 Class Imbalance Handling . . . 29
3.6 Models Configuration . . . 30
3.6.1 Experimental Environment . . . 30
3.6.2 Xception Configuration . . . 30
3.6.3 Random Forest Configuration . . . 32
3.6.4 XGBoost Configuration . . . 32
3.6.5 Final Decision — Majority Voting . . . 33
3.7 Conclusion . . . 33

4 Experimentals Results **35**
4.1 Introduction . . . 35
4.2 Evaluation Metrics . . . 35
4.3 Xception Results . . . 36
4.4 Random Forest and XGBoost Results . . . 38
4.4.1 Random Forest . . . 38
4.4.2 XGBoost . . . 39
4.4.3 Models Comparison . . . 39
4.5 XAI Results . . . 40
4.5.1 Grad-CAM Results . . . 40
4.5.2 SHAP Results . . . 42
4.5.3 Segmentation Results . . . 44
4.6 Application Presentation . . . 46
4.6.1 Analysis Page . . . 46
4.6.2 History Page . . . 48
4.6.3 Medical Report . . . 50
4.6.4 Edit Patient . . . 51
4.6.5 About Page . . . 52

4.7 Commercial Web Platform 53
4.8 Conclusion 57

5 General Conclusion **58**
5.1 Contributions 58
5.2 Limitations 59
5.3 Future Work and Perspectives 59

References **60**

List of Figures

1.1 Definition of Parkinson's Disease 2
1.2 Diagnostic Methods 3
1.3 MRI scanner and brain imaging 4
1.4 Axial, sagittal and coronal views of brain MRI 5
1.5 MRI comparison between healthy and Parkinson's disease brain 6

2.1 Relationship between Artificial Intelligence, Machine Learning and Deep Learning 8
2.2 General architecture of the proposed Parkinson's disease detection system . . . 9
2.3 General architecture of a Convolutional Neural Network 10
2.4 Difference between traditional learning and transfer learning 11
2.5 Xception architecture 13
2.6 Principle of the Random Forest algorithm 14
2.7 Principle of the XGBoost algorithm 15
2.8 Overview of Explainable AI methods 15
2.9 An example of the application of the Grad-CAM method to a brain MRI 17
2.10 SHAP Image Explainer applied to brain MRI 19
2.11 SHAP Bar plot — global feature importance across all patients 20
2.12 SHAP Beeswarm plot — SHAP value distribution across all patients 20
2.13 Segmentation of brain regions using Grad-CAM 21

3.1 General architecture of the proposed system 24
3.2 Use Case Diagram of the ParkinsonAI Application 25
3.3 Class Diagram of the ParkinsonAI Application 26
3.4 Sequence Diagram — Main Analysis Scenario 27

4.1 Confusion matrix of the Xception model 37
4.2 ROC curves of Xception, Random Forest and XGBoost 38
4.3 Accuracy comparison of Xception, Random Forest and XGBoost 40
4.4 Grad-CAM comparative heatmaps for PD and Non-PD cases 41
4.5 SHAP Image Explainer — PD vs Non-PD brain MRI regions 42
4.6 SHAP Bar plot — global clinical feature importance across all patients 43
4.7 SHAP Beeswarm plot — SHAP value distribution across all patients 43
4.8 Statistical summary — comparison between PD and Non-PD patients 44
4.9 Segmentation results based on Grad-CAM heatmaps 45
4.10 ParkinsonAI — Main Analysis page (awaiting analysis) 46
4.11 ParkinsonAI — Analysis result showing Parkinson Detected 47
4.12 ParkinsonAI — SHAP and Segmentation results 48
4.13 ParkinsonAI — Patient History page 49
4.14 ParkinsonAI — Patient detail view with model scores 49
4.15 ParkinsonAI — Patient detail view with MRI and Grad-CAM 50
4.16 ParkinsonAI — Patient detail view with SHAP and Segmentation 50
4.17 ParkinsonAI — Medical Report viewer 51
4.18 ParkinsonAI — Edit Patient popup 52
4.19 ParkinsonAI — About page 52
4.20 ParkinsonAI Web Platform — Home page 53
4.21 ParkinsonAI Web Platform — Pricing Plans 54
4.22 ParkinsonAI Web Platform — Secure Checkout page 54
4.23 ParkinsonAI Web Platform — Payment success and license delivery 55
4.24 ParkinsonAI Web Platform — Admin Dashboard 55
4.25 ParkinsonAI Web Platform — Application Interface Preview 56
4.26 ParkinsonAI Web Platform — Contact and Orders page 56

List of Tables

3.1 Dataset class distribution and applied weights 28
3.2 Data augmentation parameters applied to the training set 29
3.3 Dataset splitting ratios 29
3.4 Hardware specifications 30
3.5 Software environment 30
3.6 Xception classification head architecture 31
3.7 Xception training hyperparameters 31
3.8 Random Forest hyperparameters 32
3.9 XGBoost hyperparameters 33

4.1 Xception classification report on the test set 36
4.2 Random Forest classification report on the test set 38
4.3 XGBoost classification report on the test set 39

General Introduction

General Context

In recent years, deep learning and artificial intelligence have demonstrated remarkable potential in the field of medical image analysis, providing new opportunities for the early and accurate diagnosis of complex neurological disorders. Parkinson's disease, one of the most common and debilitating neurodegenerative diseases worldwide, affects millions of people and places a substantial burden on patients, their families, and healthcare systems.

Motivation

Early Parkinson's disease identification is still a serious clinical problem, despite tremendous advancements in medical imaging tools. Conventional diagnostic techniques may result in delayed or incorrect diagnosis since they mainly rely on clinical observation and subjective assessment. A possible way to overcome these constraints is to combine explainable AI methods with deep learning models, allowing for automated, impartial, and comprehensible interpretation of brain MRI pictures.

Objectives

The main goal of this effort is to develop and deploy an automated system that combines the transparency of explainable artificial intelligence with the capacity of deep learning to diagnose Parkinson's disease from brain MRI scans. Specifically, this study aims to:

- Train a Xception convolutional neural network for binary classification of MRI images into PD and Non-PD classes.

- Leverage the trained network as a feature extractor to feed classical machine learning models, namely Random Forest and XGBoost.
- Apply explainable AI techniques including Grad-CAM, SHAP, and segmentation to interpret and visualize the model decisions.
- Develop a desktop application integrating the complete pipeline for practical use.

Thesis Structure

This report is divided into four chapters. The first chapter provides an overview of Parkinson's disease, including its symptoms, etiology, and current diagnostic techniques. The second chapter covers the theoretical foundations of explainable artificial intelligence methods, deep learning, convolutional neural networks, and traditional machine learning algorithms. The third chapter is devoted to presenting our proposed methodology, including the proposed architecture and its main modules. Finally, the fourth chapter presents the developed desktop application and discusses the experimental results, including the XAI analysis and the performance evaluation of all the proposed models.

Chapter 1

Parkinson's Disease and Diagnostic Methods

1.1 Introduction

Parkinson's disease, also called Parkinson's syndrome, is a progressive neurological condition that affects a person's ability to control their movements. It is thought to be one of the most prevalent neurological conditions in the world, especially affecting the elderly. It is caused by the death of nerve cells in a particular area of the brain called the substantia nigra, which results in a deficiency of dopamine, a neurotransmitter essential for motor function.

Since no single test can conclusively demonstrate the existence of this disease, early identification is one of the biggest obstacles. Nonetheless, magnetic resonance imaging (MRI) is essential for detecting structural alterations in the brain, which makes it easier to identify the illness early on.

The goal of this chapter is to give a thorough explanation of Parkinson's disease, including how it can be diagnosed, how MRI can be used to detect it, and how artificial intelligence can help with diagnosis.

1.2 Definition and Impact of Parkinson's Disease

Parkinson's disease is the second most common neurological disorder, affecting more than 10 million people worldwide. Its prevalence increases to a higher level after 60 years of age [1]. Every year, the disease costs billions of dollars around the world, putting a heavy strain on health services and the families of people who suffer from it.



Figure 1.1: Definition of Parkinson's Disease

The loss of dopaminergic neurons in the substantia nigra causes Parkinson's disease (PD), a progressive neurodegenerative disorder that impairs motor function and lowers dopamine levels [2]. Bradykinesia, postural instability, muscle rigidity, and resting tremor are the main motor signs. Additionally, frequently noted non-motor symptoms include sadness, sleep issues, and cognitive decline [3].

1.3 Diagnostic Methods of Parkinson's Disease

Parkinson's disease diagnosis remains a clinical challenge because there isn't a single, accurate test to detect the illness. The diagnosis is mostly based on a neurological examination, a medical history, and the observation of characteristic motor signs [4].

Many diagnostic techniques are commonly used. Neuroimaging techniques, such as Magnetic Resonance Imaging (MRI) and Dopamine Transporter Scanning (DaTSCAN), help detect structural and functional changes in the brain associated with the disease [5]. Blood testing and cerebrospinal fluid studies may also be used to rule out other conditions.

By automatically evaluating clinical data and medical imaging, recent advances in artificial intelligence have made it feasible to detect Parkinson's disease more accurately and early [6].

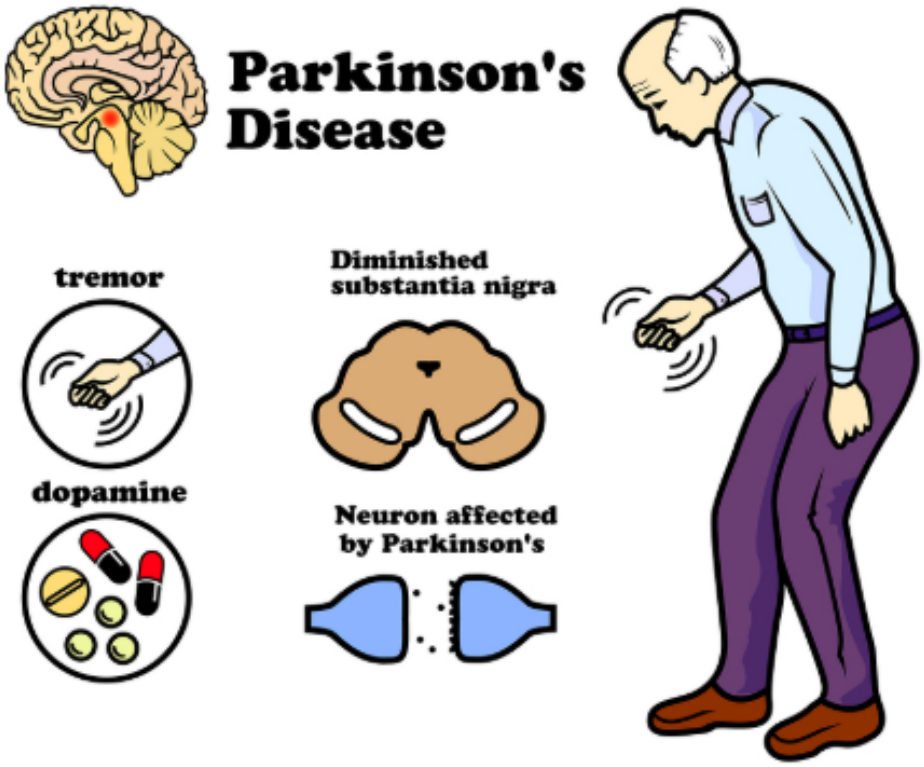


Figure 1.2: Diagnostic Methods

These technologies provide a practical way to improve diagnostic accuracy and reduce the rate of misdiagnosis.

In this study, we apply artificial intelligence techniques, more precisely deep learning, to analyze MRI images of the brain and detect whether the patient suffers from Parkinson's disease. In the next section, we present an overview of the magnetic resonance imaging (MRI) system.

1.4 Magnetic resonance imaging (MRI)

1.4.1 Definition of MRI

Magnetic resonance imaging (MRI) is a non-invasive medical imaging method that produces detailed images of the body's internal organs. MRI is a safer alternative for frequent use in clinical settings because it does not use ionizing radiation like CT or X-rays do [7]. Because it creates high-contrast images of soft tissues, it is particularly useful for studying the brain.

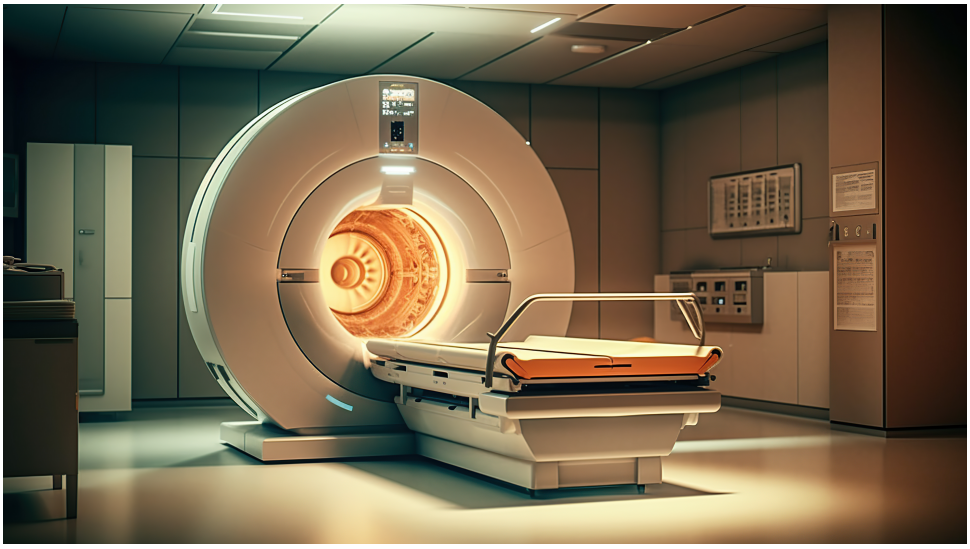


Figure 1.3: MRI scanner and brain imaging

1.4.2 Principles and Basic Concepts of MRI

MRI takes advantage of the magnetic properties of hydrogen nuclei present in biological tissues. When a patient is subjected to a strong magnetic field, these nuclei align along the field axis. A radiofrequency pulse is then applied to disrupt this alignment. When the nuclei return to equilibrium, they produce signals that are captured and converted into complex images [7].

The main MRI concepts include:

- **T1 and T2 contrasts:** Two types of MRI contrasts used to analyze different tissue structures.
- **Diffusion MRI (DTI):** A technique for studying neural connectivity by measuring water molecule displacement.
- **Functional MRI (fMRI):** Allows visualization of brain activity by measuring blood oxygenation variations.

Figure 1.4 presents the different views of the brain obtained from the MRI device.

1.4.3 MRI for the Diagnosis of Parkinson's Disease Symptoms

With the use of MRI, it is possible to detect and evaluate the symptoms and brain alterations associated with Parkinson's disease. MRI is essential for diagnosing and monitoring Parkinson's disease because it can detect structural and functional changes in the brain associated

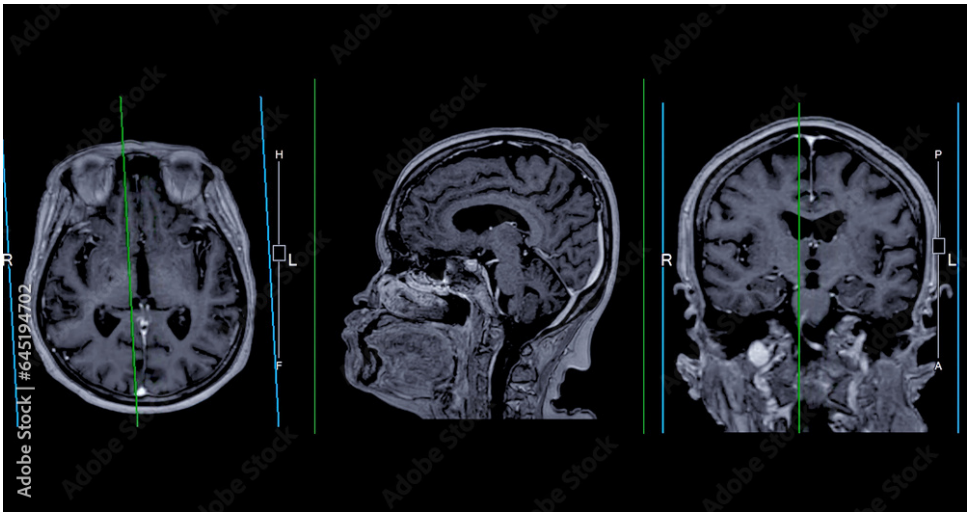


Figure 1.4: Axial, sagittal and coronal views of brain MRI

with neurodegeneration [8].

1.4.3.1 Substantia Nigra Abnormalities

One of the initial symptoms of Parkinson's disease is substantia nigra degeneration. Thanks to advanced MRI techniques like neuromelanin-sensitive imaging, this region may be observed and dopaminergic neuron loss can be assessed.

1.4.3.2 White Matter Alterations

Diffusion tensor imaging (DTI) can be used to evaluate white matter integrity, which is essential for brain signal transmission. Parkinson's disease patients frequently show signs of these connections deteriorating.

1.4.3.3 Functional Brain Changes

Functional magnetic resonance imaging (fMRI) makes it possible to examine brain activity by identifying blood flow in different parts of the brain. Parkinson's patients often have decreased activity in the motor cortex and basal ganglia [8].

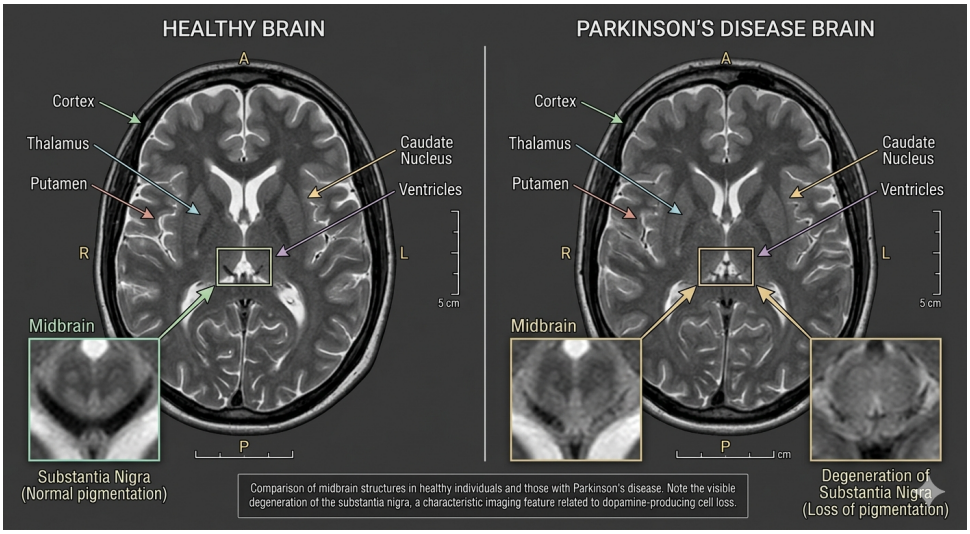


Figure 1.5: MRI comparison between healthy and Parkinson's disease brain

1.5 The Role of Artificial Intelligence in Diagnosis

Through automated analysis of clinical data and medical pictures, artificial intelligence (AI) has greatly increased the diagnosis of Parkinson's disease, allowing for more accurate and early detection of this disease[9].

Convolutional Neural Networks (CNNs) are one of the best methods for analyzing MRI data to find structural brain changes associated with Parkinson's disease [10]. These algorithms automatically extract intricate information from photos, outperforming traditional hand analysis methods. Automated brain structure segmentation, illness severity categorization, and disease progression prediction are important AI applications in Parkinson's disease diagnosis. Despite these advancements, a number of obstacles still stand in the way of broader use, such as the lack of annotated datasets, model interpretability, and clinical integration [9].

1.6 Conclusion

An overview of Parkinson's disease, including its symptoms and effects on patients, was given in this chapter. Next, we looked at the many early diagnostic techniques, focusing especially on the function of magnetic resonance imaging as a crucial instrument for identifying structural alterations in the brain. Lastly, we looked at how artificial intelligence might help with early disease identification and increased diagnostic precision.

Chapter 2

Deep Learning Models and Machine Learning Algorithms for Parkinson's Detection

2.1 Introduction

In this chapter, we are presented The models and methods utilized for MRI image classification in the context of Parkinson's disease detection . Convolutional Neural Networks (CNNs) based on deep learning and conventional machine learning classifiers are the two primary families of techniques investigated. Additionally, we discuss the Explainable AI (XAI) techniques that are employed to analyze our models' predictions.

2.2 Machine Learning, Deep Learning and Supervised Classification

2.2.1 Machine Learning

At the intersection of computer science, data science, and artificial intelligence, machine learning is a field that allows computers to automatically learn from data and improve without the need for explicit programming for every task [11]. Machine learning algorithms au-

tonomously learn to perform a task or make predictions from data and improve their performance over time. Once trained, the algorithm is able to identify patterns in new data.

Machine learning paradigms can be mainly divided into three categories: supervised learning, unsupervised learning, and reinforcement learning [11]. Supervised classification is a supervised learning task in which a model learns to assign a class (label) to a piece of data based on already labeled examples. More precisely, every training sample has input features that reflect the properties of the data, together with a label that indicates which class those features belong to. The classification algorithm aims to build a model that can predict the unknown class of new data based on its attributes [12].

2.2.2 Deep Learning

Inspired by the structure of the human brain, deep learning is a branch of machine learning that employs multi-layered artificial neural networks [12]. These networks stand out for their capacity to handle data in a hierarchical and flexible way, acquiring progressively more abstract representations from unprocessed input. Figure 2.1 shows the connection between deep learning, machine learning, and artificial intelligence.

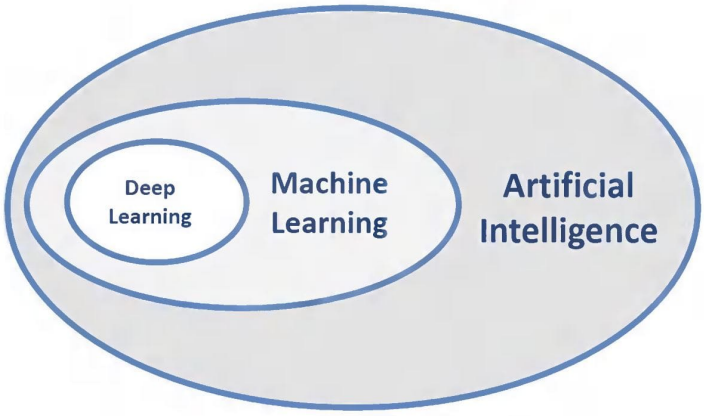


Figure 2.1: Relationship between Artificial Intelligence, Machine Learning and Deep Learning

2.2.3 Advantages of Deep Learning over Machine Learning

The process of extracting features from data is the primary benefit of deep learning over traditional machine learning. The quality of these characteristics significantly affects the model's performance in traditional machine learning approaches, which frequently require the human

acquisition of pertinent information utilizing outside techniques or domain knowledge. On the other hand, deep learning extracts implicit and automated features straight from raw data [11]. Convolutional Neural Networks (CNNs) are especially helpful for processing unstructured data like MRI images because their convolutional layers gradually learn more abstract and discriminative representations, doing away with the need for an explicit feature extraction phase [12].

2.2.4 Architecture of a Computer-Aided Diagnosis System for Parkinson's Disease

In this work, we automatically identify Parkinson's disease from brain MRI scans using a supervised classification approach. The proposed method is based on a hybrid architecture that combines traditional machine learning techniques with deep learning techniques. The Xception CNN automatically extracts features from MRI images, while the final supervised classification is performed using Random Forest and XGBoost classifiers. Moreover, the proposed architecture incorporates Explainable Artificial Intelligence (XAI) techniques, such as Grad-CAM, SHAP, and segmentation, to provide comprehensible justifications for the model's predictions. The overall architecture of the proposed system is depicted in Figure 2.2.

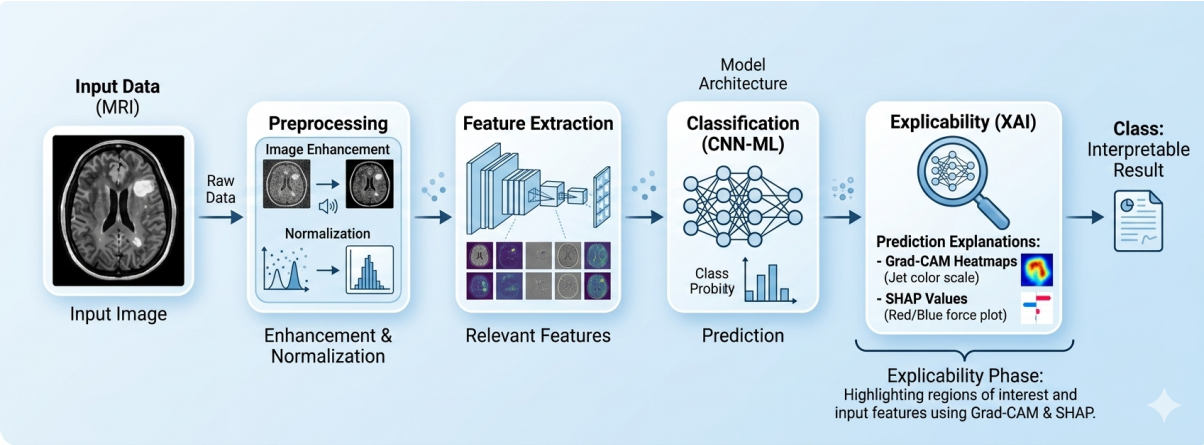


Figure 2.2: General architecture of the proposed Parkinson's disease detection system

2.3 Convolutional Neural Networks (CNNs)

A family of deep learning models called Convolutional Neural Networks (CNNs) was created especially to handle grid-structured data, including pictures [12]. They have shown exceptional performance in tasks including object detection, segmentation, and picture classification,

especially in medical image analysis. CNNs automatically learn features using convolutional layers, which identify visual patterns like edges, textures, and forms, in contrast to traditional neural networks that need manual feature extraction [11]. As shown in Figure 2.3, a typical Convolutional Neural Network (CNN) architecture is composed of two main blocks: a feature extraction block and a classification block.

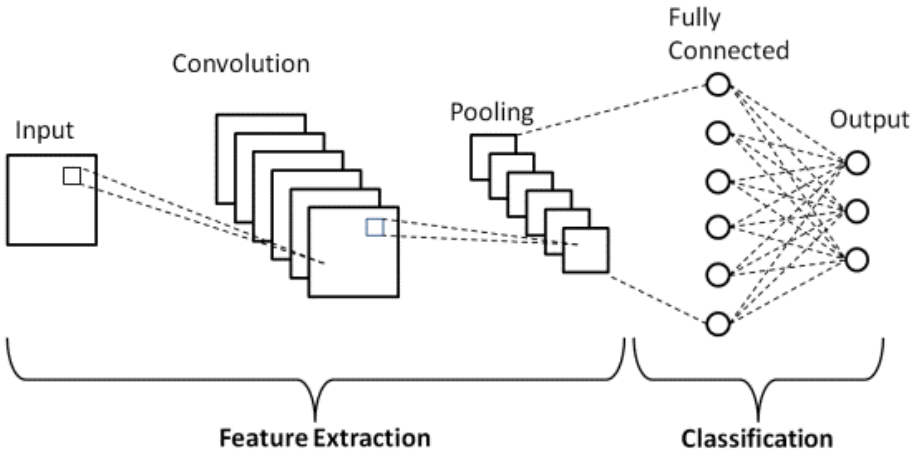


Figure 2.3: General architecture of a Convolutional Neural Network

2.3.1 Feature Extraction Block

The central part of a CNN is the feature extraction block. Convolutional filters are applied to the input image to produce feature maps, which are then followed by activation functions. The hierarchical extraction of visual information is made possible by repeating this process across several layers. Convolutional layers, which identify spatial patterns, and pooling layers, which decrease the spatial dimensions of feature maps while preserving crucial information, make up the block's two primary layer types.

2.3.2 Classification Block

The fully connected layers that make up the classification block are trained to link the retrieved feature vector to a particular class. The three-dimensional feature maps are first converted into a one-dimensional vector via a flatten layer. The output layer then calculates the likelihood of belonging to each class using a softmax activation function [11].

CNNs have several hyperparameters, including the learning rate, number of epochs, filter size, stride, and padding. These parameters play an important role in controlling the training

process and the architecture of the network. Padding helps preserve the spatial dimensions of the feature maps, while the stride determines the step size of the filter during the convolution operation. The ability of the model to learn and extract meaningful representations from the training data is directly influenced by these hyperparameters.

2.4 Transfer Learning and Pre-trained Models

Transfer learning refers to methods that transfer knowledge acquired from one problem to solve another. This approach has made significant progress with the advent of deep learning, as deep models often require substantial resources and long training times. By using pre-trained models as a foundation, transfer learning accelerates the development of high-performance models for complex tasks, particularly in computer vision and natural language processing (NLP) [13]. Figure 2.4 illustrates the concept of transfer learning, which allows us to take advantage of knowledge already acquired from massive datasets such as ImageNet, which contains millions of annotated images, rather than training a model from scratch.

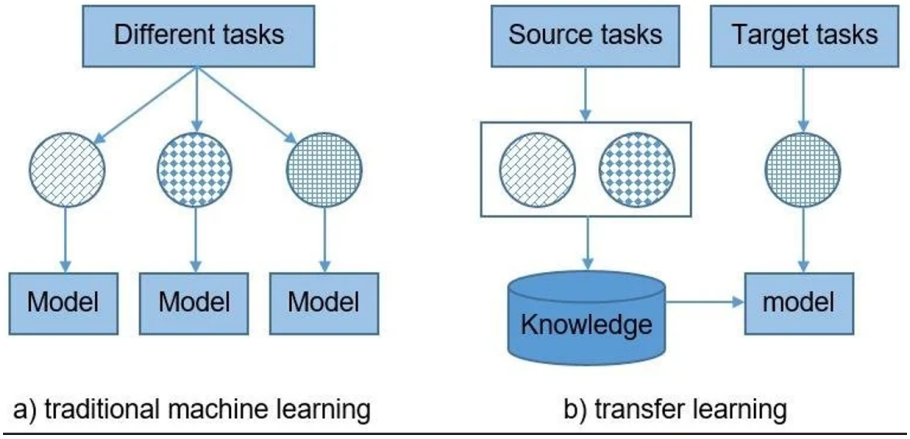


Figure 2.4: Difference between traditional learning and transfer learning

Before discussing transfer learning strategies, we first present its fundamental concepts, including pre-trained models and the key advantages of transfer learning.

2.5 Transfer Learning Strategies

Depending on the size and similarity of the target dataset, there are three primary approaches for using transfer learning.

2.5.1 Strategy 1: Full Training

In the full training strategy, all layers are retrained using the new dataset, with the pre-trained model serving merely as an initialization point. This approach is appropriate when the destination dataset is sizable and sufficiently distinct from the source dataset.

2.5.2 Strategy 2: Partial Training

CNNs learn hierarchical representations: the initial layers capture general features (e.g., edges and textures), while deeper layers extract task-specific features [13]. This strategy consists of freezing the initial layers to preserve these general representations and retraining only the deeper layers to adapt the model to the new task. Note that this method is appropriate when the target dataset is sufficiently large and relatively similar to the source dataset.

2.5.3 Strategy 3: Freezing all Convolutional Layers

In this strategy, only the classification head is retrained, and all convolutional layers are retained. This approach can be adopted when the target dataset is small and comparable to the source dataset [14].

2.6 Transfer Learning Process

2.6.1 Pre-trained Model Selection

There exist a variety of pre-trained models, The type of the task at hand determine which pre-trained model is best. In our study, we select Xception because of its effective use of depthwise separable convolutions and its good performance on image classification challenges.

2.7 Xception

Chollet [15] proposed the deep convolutional neural network architecture known as Xception, or Extreme Inception. It is predicated on the idea that it is possible to completely separate the mapping of spatial and cross-channel correlations in the feature maps of convolutional neural networks.

2.7.1 Xception Architecture

Xception is an extension of the Inception architecture in which the classic modules of Inception are replaced by depthwise separable convolutions (a depthwise convolution followed by a pointwise convolution) [13]. This method maintains high classification performance while drastically reducing the size of the model [15].

The architecture of Xception is depicted in Figure 2.5, it can be seen that, Xception is composed of 36 convolutional layers, which are divided into three main flows: the Entry flow, the Middle flow, which is repeated eight times, and the Exit flow. Similar to ResNet, the model takes advantage of residual connections to help avoid the vanishing gradient problem during the training process.

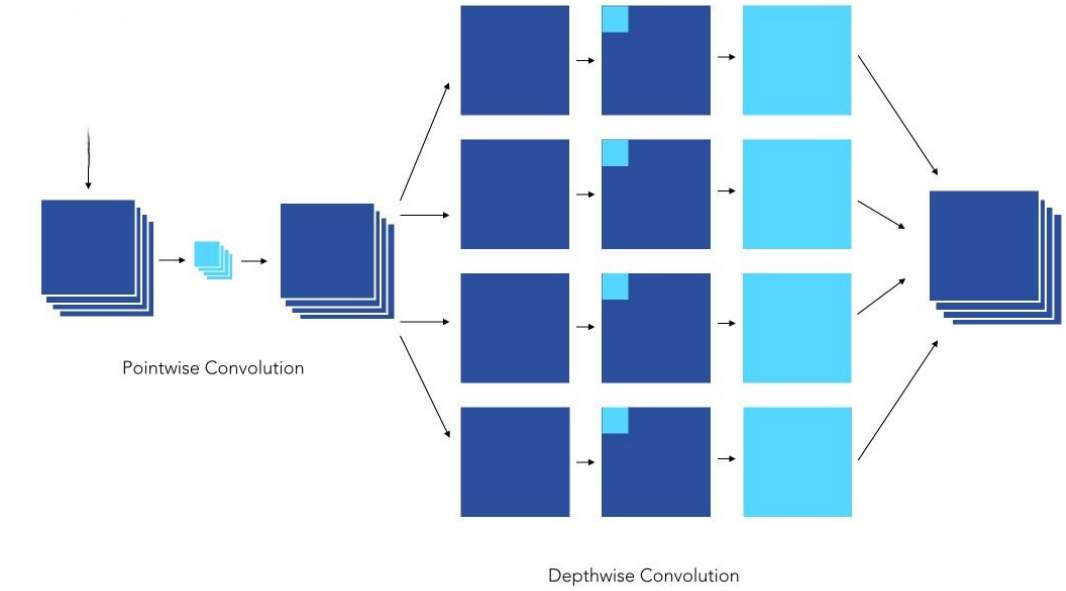


Figure 2.5: Xception architecture

The main feature extractor used in our work is the Xception model, which is pre-trained on ImageNet. The GlobalAveragePooling2D layer generates a 2048-dimensional feature vector that is then fed into the classification block for Parkinson's disease prediction. The next section presents the main classification methods employed in our study.

2.8 Classification Methods

In this work, two machine learning classifiers are adopted to perform the final categorization of Parkinson's disease from MRI images after extracting features using Xception. The two classifiers are: Random Forest and XGBoost.

2.8.1 Random Forest (RF)

During training, Random Forest, an ensemble learning technique, builds several decision trees and outputs the class that is the mode of the classes of the individual trees [16]. In order to minimize overfitting and enhance generalization, each tree is trained using a random subset of the training data and a random selection of features.

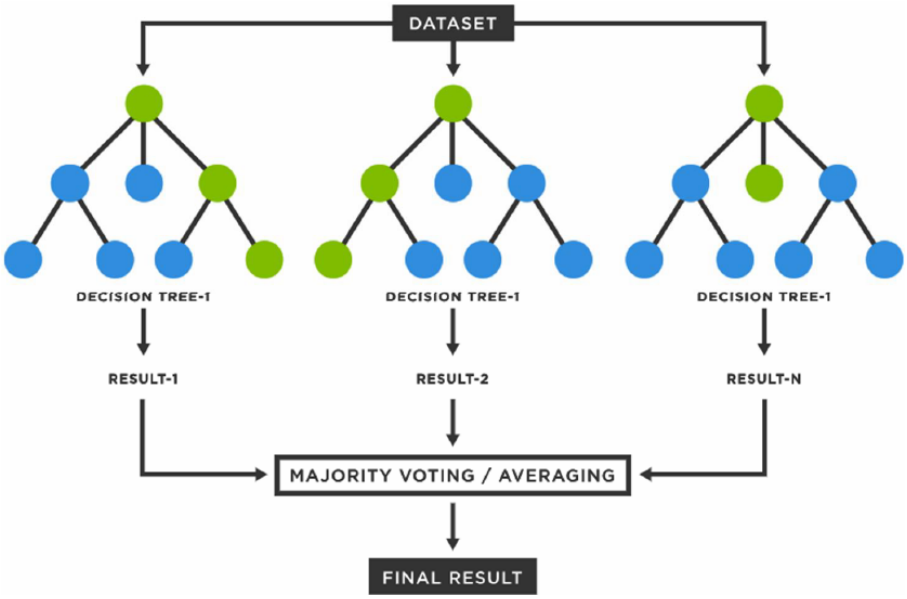


Figure 2.6: Principle of the Random Forest algorithm

2.8.2 XGBoost

A scalable and effective gradient boosting implementation is called XGBoost (Extreme Gradient Boosting) [17]. It constructs a group of weak learners in a sequential manner, with each new tree fixing the mistakes of the preceding ones. XGBoost is renowned for its quick computation and excellent performance.

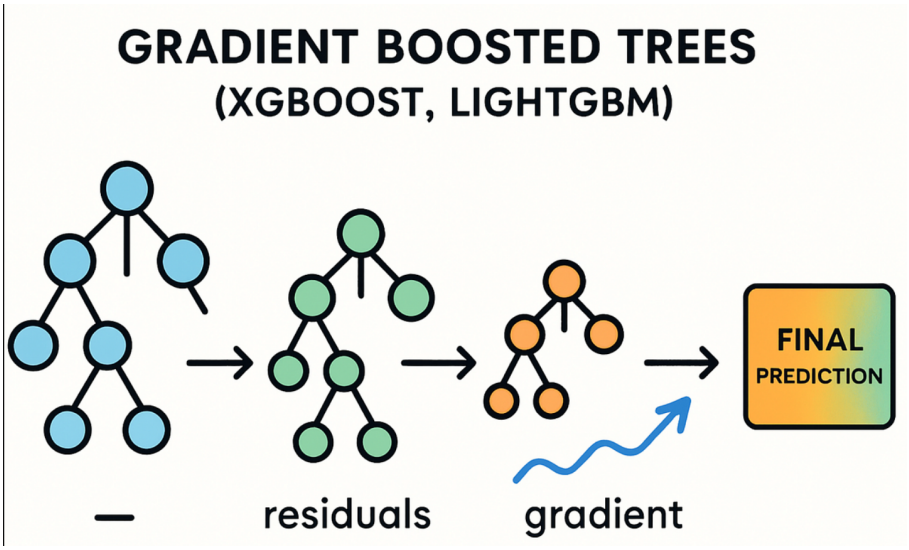


Figure 2.7: Principle of the XGBoost algorithm

2.9 Explainable AI (XAI)

The term "explainable artificial intelligence" (XAI) describes strategies and tactics that enable humans to comprehend and interpret the predictions of AI models [18]. Interpretability is essential in medical applications because it enables physicians to comprehend and have faith in the choices made by AI systems.

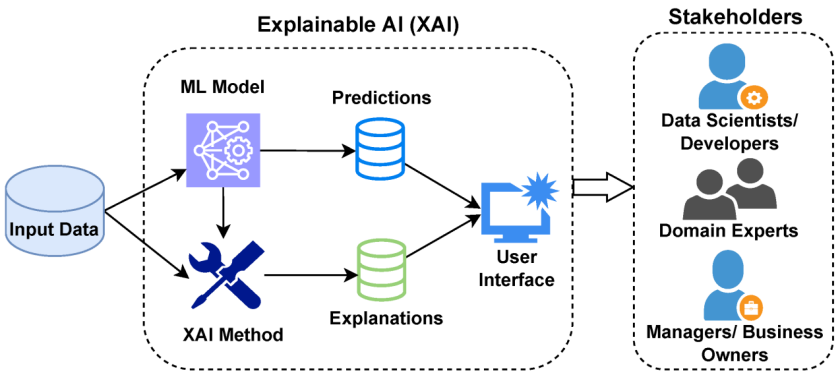


Figure 2.8: Overview of Explainable AI methods

Three XAI techniques—Grad-CAM for visual explanation, SHAP for feature importance analysis, and unsupervised segmentation for finding aberrant brain regions—are used in our work to interpret the predictions of our Xception model.

2.9.1 Grad-CAM

Gradient-weighted Class Activation Mapping (Grad-CAM) is a visual explanation technique that uses the gradients of the target class flowing into the final convolutional layer to build a coarse localization map that highlights the important portions of the image for class prediction [19]. This technique belongs to the broader family of saliency-based explanation methods, which aim to identify the input regions that contribute most to a deep network's decision [18]. Grad-CAM's ability to be applied directly to an already-trained CNN, like the Xception model employed in this study, without requiring any architectural alteration or retraining is one of its primary advantages [19].

2.9.1.1 General Functioning of Grad-CAM

The construction of a Grad-CAM heatmap relies on a sequence of simple steps that transform the internal information of the network into an explanatory visual map [19]:

- **Target score computation:** After processing the MRI image, the model generates a score for each class (PD/Non-PD); the score that corresponds to the class that needs to be explained is chosen.
- **Gradient computation:** The final convolutional layer's feature maps are used to calculate the gradient of this target score, which shows which areas have the greatest influence on the choice.
- **Channel weighting:** An important weight α_k^c is obtained by averaging the gradients for each channel. The related channel's significance for the prediction increases with this weight.
- **Heatmap generation:** After multiplying each feature map by its corresponding weight, the results are summed. Next, only the positive impacts are retained by applying a ReLU function. Lastly, the final map is adjusted to fit the original image's proportions.

2.9.1.2 Mathematical Formulation of Grad-CAM

There are two basic processes in the Grad-CAM map computation [19].

1. Feature importance (α_k^c): The gradient of the raw score y^c is calculated with respect to

these activations and spatially averaged using Global Average Pooling [20] in order to assess the sensitivity of each feature map A^k of the final convolutional layer with regard to the predicted class c (e.g., Parkinson's disease):

$$\alpha_k^c = \frac{1}{Z} \sum_i \sum_j \frac{\partial y^c}{\partial A_{i,j}^k} \tag{2.1}$$

where Z is the feature map's entire pixel count (for example, 14×14). This is equivalent to using `tf.GradientTape()` and `tf.reduce_mean()` in Python.

2. Linear combination and ReLU filtering: The weights α_k^c are then linearly combined with the maps A^k. The application of the ReLU function allows the removal of negative gradients (noise) in order to retain only the visual elements that have a positive influence on the diagnosis:

$$L_{\text{Grad-CAM}}^c = \text{ReLU}\left(\sum_k \alpha_k^c A^k\right) \tag{2.2}$$

In code, this operation is materialized by the line: `heatmap = np.maximum(heatmap, 0)`.

In our work, Grad-CAM provides visual proof of the areas impacted by Parkinson's disease by highlighting the brain regions that most influenced the model's prediction, particularly the Substantia Nigra and Basal Ganglia regions known to be affected by the disease [2]. Figure 4.4 shows an example of the application of the Grad-CAM method to a brain MRI.

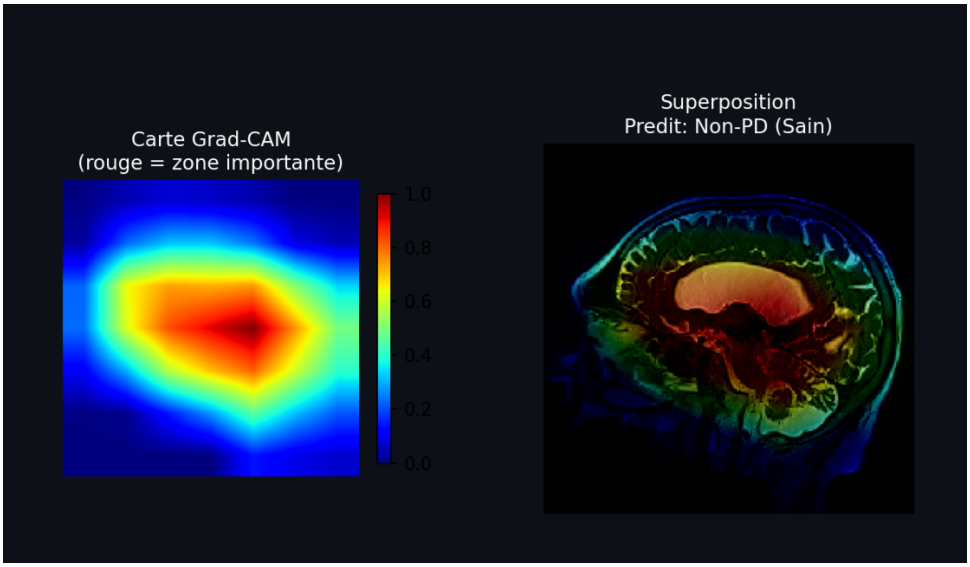


Figure 2.9: An example of the application of the Grad-CAM method to a brain MRI

2.9.2 SHAP

SHAP (SHapley Additive exPlanations) [21] is a cooperative game theory-based unified framework for assessing model predictions. The main concept is derived from the Shapley value, which was first presented by Lloyd Shapley in 1953 [22]. In this model, each feature is viewed as a "player" in a coalition, and the model's prediction for a particular sample is viewed as the payout to be divided among the features.

Formally, the Shapley value ϕ_i of a feature i is given by:

$$\phi_i(v) = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!\,(|N| - |S| - 1)!}{|N|!} \Big(v(S \cup \{i\}) - v(S) \Big) \tag{2.3}$$

where $v(S)$ is the value function that represents the model's predicted output when just the features in S are known, N is the set of all features, and S is a coalition that does not contain i [21]. Additivity is a crucial aspect of this formulation: for a sample x, the model's prediction $f(x)$ may be written as the sum of the SHAP values of all its features and a baseline (average) prediction ϕ_0:

$$f(x) = \phi_0 + \sum_{j=1}^{M} \phi_j \tag{2.4}$$

The relevance value ϕ_j given to each attribute for a given prediction indicates how much each item influences the final conclusion. In this study, SHAP is applied in two complementary ways: first, as a feature significance analyzer on clinical data using a Random Forest classifier, and second, as an image explanation directly on the MRI images processed by the Xception CNN.

2.9.2.1 SHAP Image Explainer

Since it necessitates analyzing every $2^{|N|}$ possible coalition, the exact computation of Equation (2.3) is typically prohibitively expensive [21]. SHAP values are estimated using a **Partition Explainer** in conjunction with a *blur masker*, a model-agnostic, perturbation-based method tailored to image data [21], since the "features" for the Xception CNN operating on MRI images correspond to image regions rather than a small set of tabular variables.

The SHAP Image Explainer examines how various image areas contribute to the model's prediction using this blur masker. A heatmap is created for every MRI picture, with blue re-

gions pushing against a certain prediction (PD or non-PD) and red/pink regions representing locations that pushed the model in that direction. The influence on the model's output increases with color intensity, consistently with the additivity property of Equation (2.4), where ϕ_0 corresponds to the model's average output over the dataset.

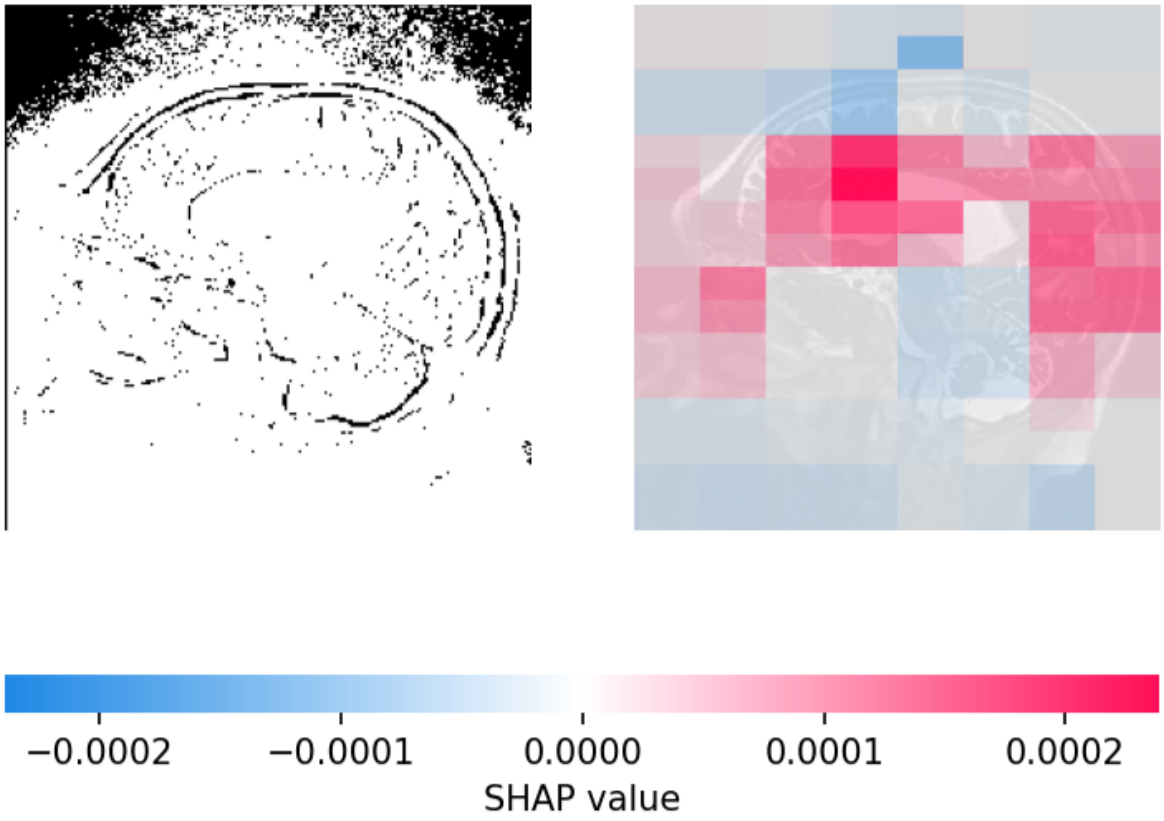


Figure 2.10: SHAP Image Explainer applied to brain MRI

2.9.2.2 SHAP on Clinical Data

The second approach involves applying SHAP to a Random Forest classifier [16] that was trained using clinical data from the NTUA Parkinson dataset CSV files [23]. The seven clinical characteristics used are age, UPDRS I, UPDRS II, UPDRS III, MMSE, FOG, and sex. Using a `TreeExplainer` [21], which takes advantage of the Random Forest's inherent tree structure rather than directly evaluating Equation (2.3), SHAP values for tree-based models are calculated *exactly* and efficiently (in polynomial rather than exponential time). With ϕ_0 representing the average predicted probability of Parkinson's disease over the training set, the resulting SHAP values for each patient meet the additivity property of Equation (2.4).

To understand the model's choices for every patient, two global visuals are created:

- **SHAP Bar plot:** displays the mean absolute SHAP value for each clinical feature across all patients, ranking attributes based on their overall influence on the model's conclusions.
- **SHAP Beeswarm plot:** shows the distribution of SHAP values for each patient's characteristics. Each dot represents a single patient; red dots indicate high feature values and blue dots indicate low feature values. Dots on the right represent PD prediction, whereas dots on the left represent non-PD prediction.

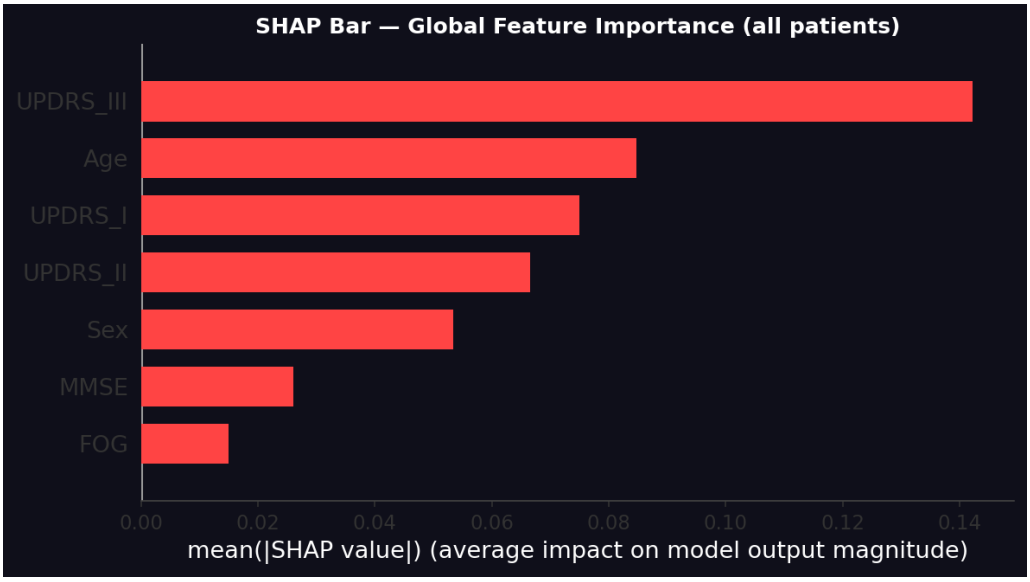


Figure 2.11: SHAP Bar plot — global feature importance across all patients

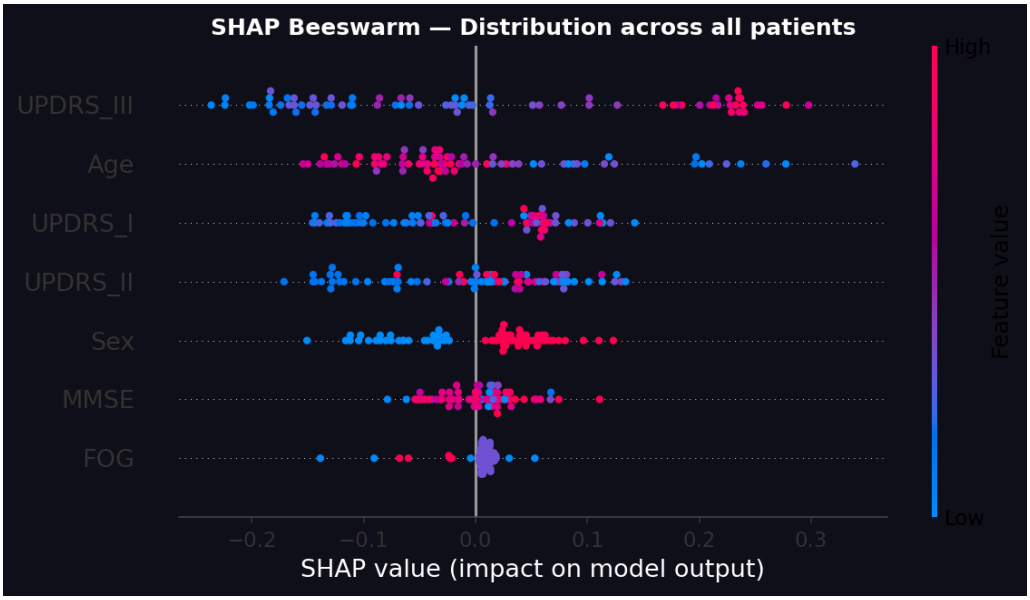


Figure 2.12: SHAP Beeswarm plot — SHAP value distribution across all patients

2.9.3 Segmentation

Unsupervised segmentation using Grad-CAM heatmaps can detect abnormal brain regions without the requirement for human annotations [19]. To highlight regions of interest, a binary mask is created using the Grad-CAM activation map derived from Equation (2.2). Formally, the binary segmentation mask $M(i,j)$ is obtained by thresholding the normalized Grad-CAM map $L^c_{\text{Grad-CAM}}$ with a threshold $\tau \in [0,1]$:

$$M(i,j) = \begin{cases} 1, & \text{if } L^c_{\text{Grad-CAM}}(i,j) \geq \tau \\ 0, & \text{otherwise} \end{cases} \tag{2.5}$$

where pixels below this threshold are removed and pixels with an activation value above τ are regarded as belonging to an aberrant region. Instead of producing a graded picture, this procedure transforms the continuous heatmap into a binary demarcation of the regions of interest, in line with classical thresholding approaches used in medical image segmentation [24]. The segmentation mask is empty in non-PD cases, indicating a normal brain; in PD cases, it displays the abnormal regions in red, providing a visual representation of the disease's effects on the brain, particularly around the Substantia Nigra and Basal Ganglia [2].

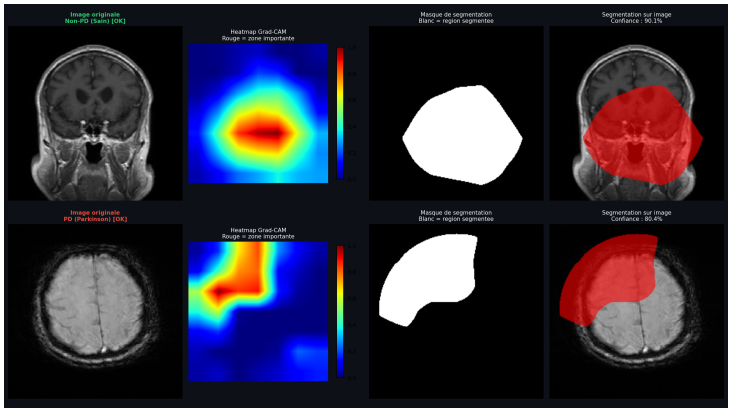


Figure 2.13: Segmentation of brain regions using Grad-CAM

2.10 Conclusion

This chapter covered the fundamental ideas of deep learning, convolutional neural networks, and transfer learning in order to present the theoretical underpinnings of the suggested system. Random Forest and XGBoost were presented as complementary classifiers, and the

Xception architecture was emphasized as the main feature extractor. Lastly, Grad-CAM and SHAP, two XAI approaches, were discussed as ways to guarantee clinical interpretability and transparency. Together, these components provide the methodological foundation for the experimental work that is described in the next chapter.

Chapter 3

Proposed Methodology

3.1 Introduction

This chapter describes the practical approach used to identify Parkinson's disease from brain MRI images, building on the theoretical underpinnings laid out in the preceding chapter. The goal is to thoroughly explain every stage of the pipeline that has been put into place, from model configuration to data preparation. Our method is divided into two main stages: the Xception convolutional neural network is used directly as a classifier in the first phase, and its output is used as a feature extractor in the second phase, feeding Random Forest and XGBoost, two traditional machine learning models. The presentation of the dataset, the image preprocessing procedures, and the setup of every model utilized in this investigation are all covered in this chapter.

3.2 The Proposed System's General Architecture

The general architecture of the suggested system is shown in Figure 3.1. The NTUA Parkinson dataset [23] is where the pipeline starts, including MRI pictures from people with Parkinson's disease and those without the disease. Before the images are sent into the Xception CNN, they go through a preprocessing step [15]. Xception uses straight binary categorization in the first stage. In the second stage, each image's 2048 features are extracted by Xception's GlobalAveragePooling layer. These features are then reduced to 500 using SelectKBest [25], then Random Forest [16] and XGBoost [17] are used for classification. To ensure that the findings are interpretable, XAI techniques like Grad-CAM [19], SHAP [21], and segmentation are added to the entire system.

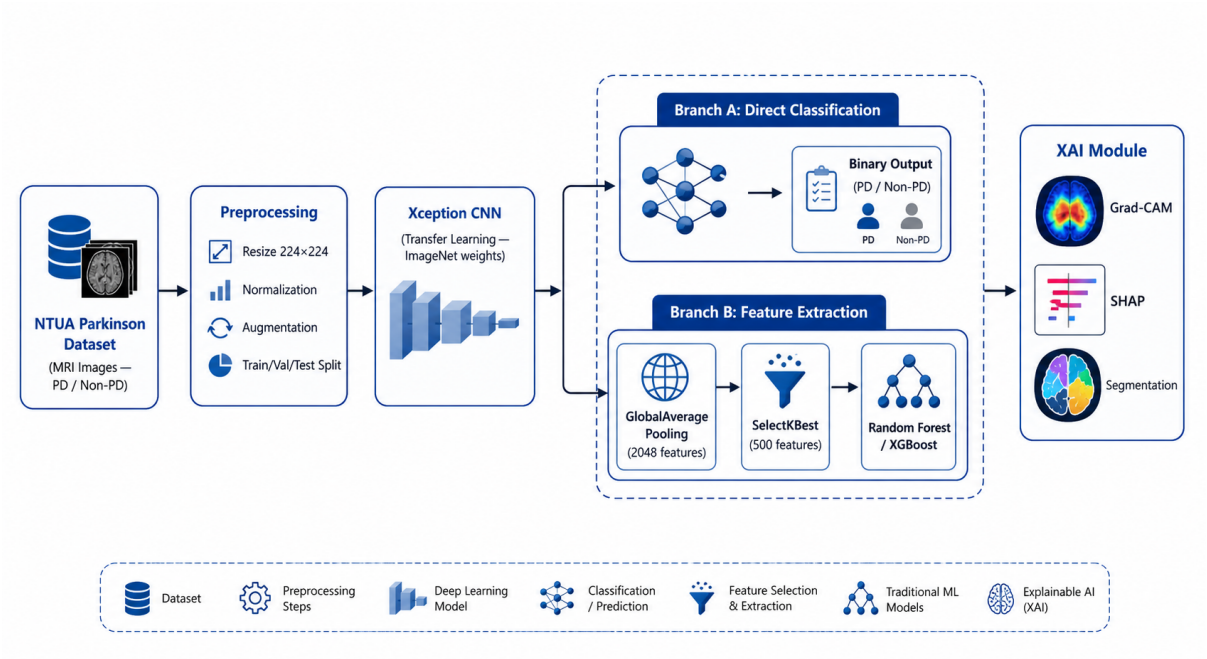


Figure 3.1: General architecture of the proposed system

3.3 UML Diagrams

3.3.1 Use Case Diagram

Figure 3.2 presents the use case diagram of the ParkinsonAI desktop application. The system identifies a single actor, the **User** (medical professional), who interacts with the application through three main modules: image analysis, patient history management, and system information.

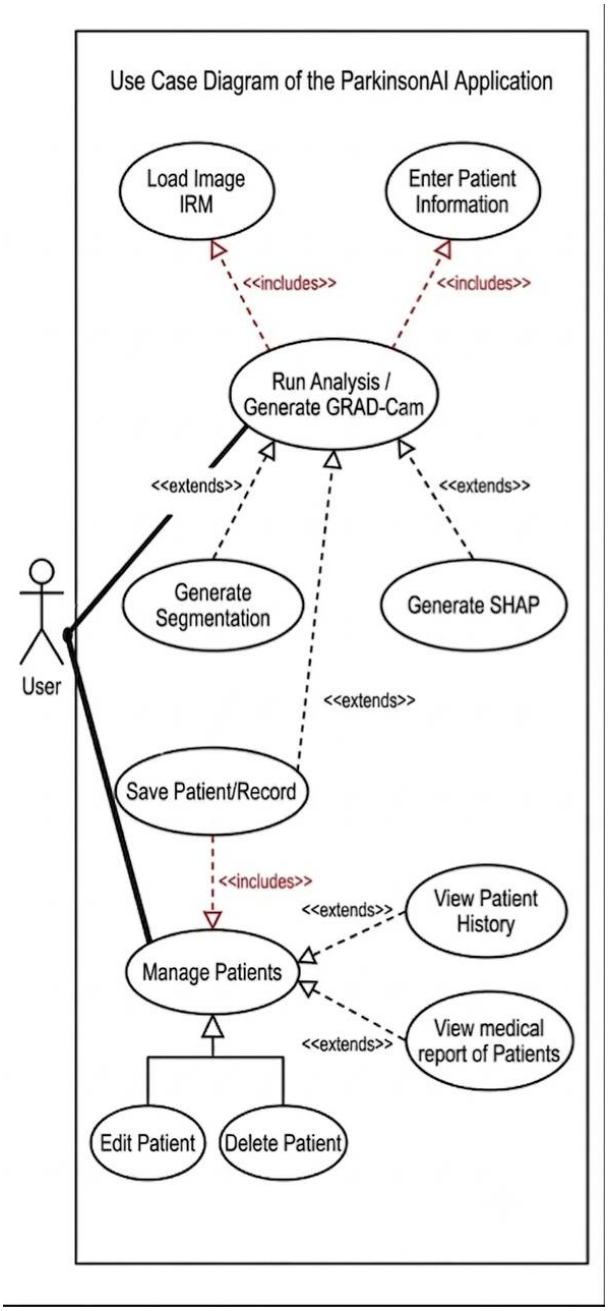


Figure 3.2: Use Case Diagram of the ParkinsonAI Application

3.3.2 Class Diagram

Figure 3.3 presents the class diagram of the ParkinsonAI application, derived from the conceptual data model. Each entity is mapped to a class with its corresponding attributes and methods.

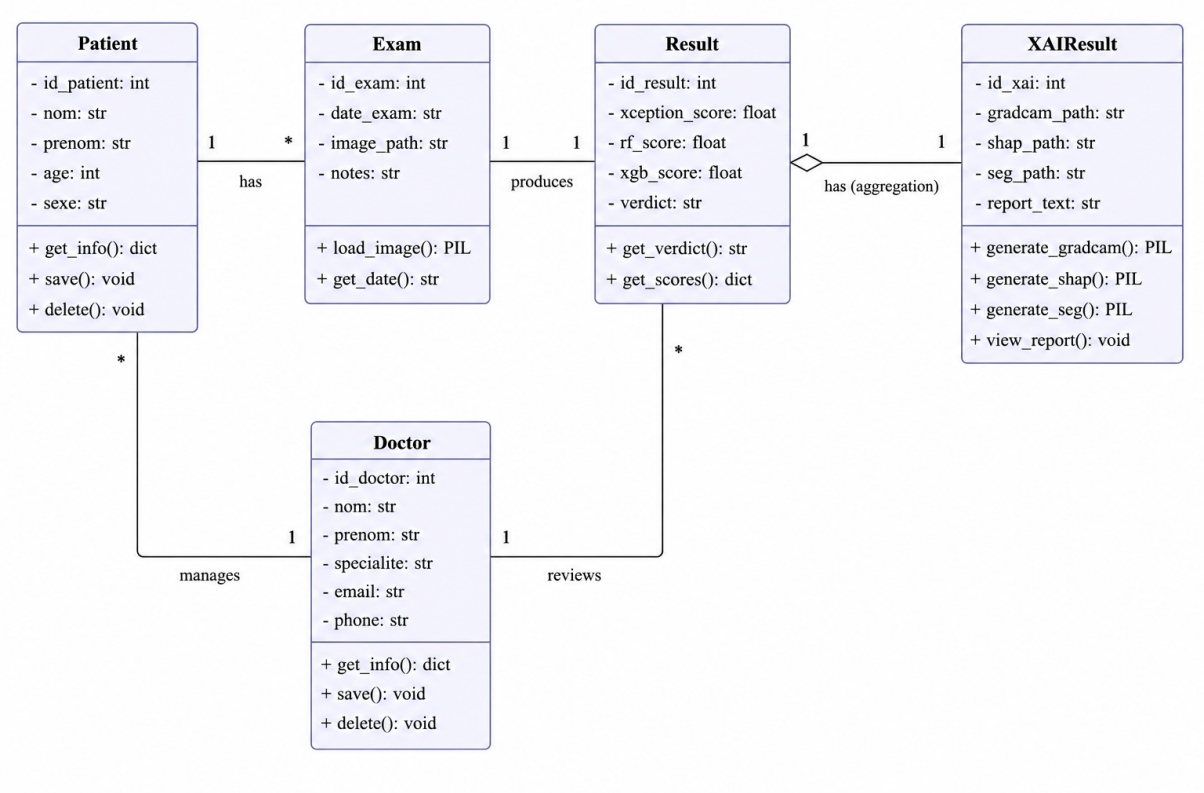


Figure 3.3: Class Diagram of the ParkinsonAI Application

3.3.3 Sequence Diagram

Figure 3.4 illustrates the sequence diagram of the main analysis scenario in the ParkinsonAI application, from image loading to saving the patient record in the database.

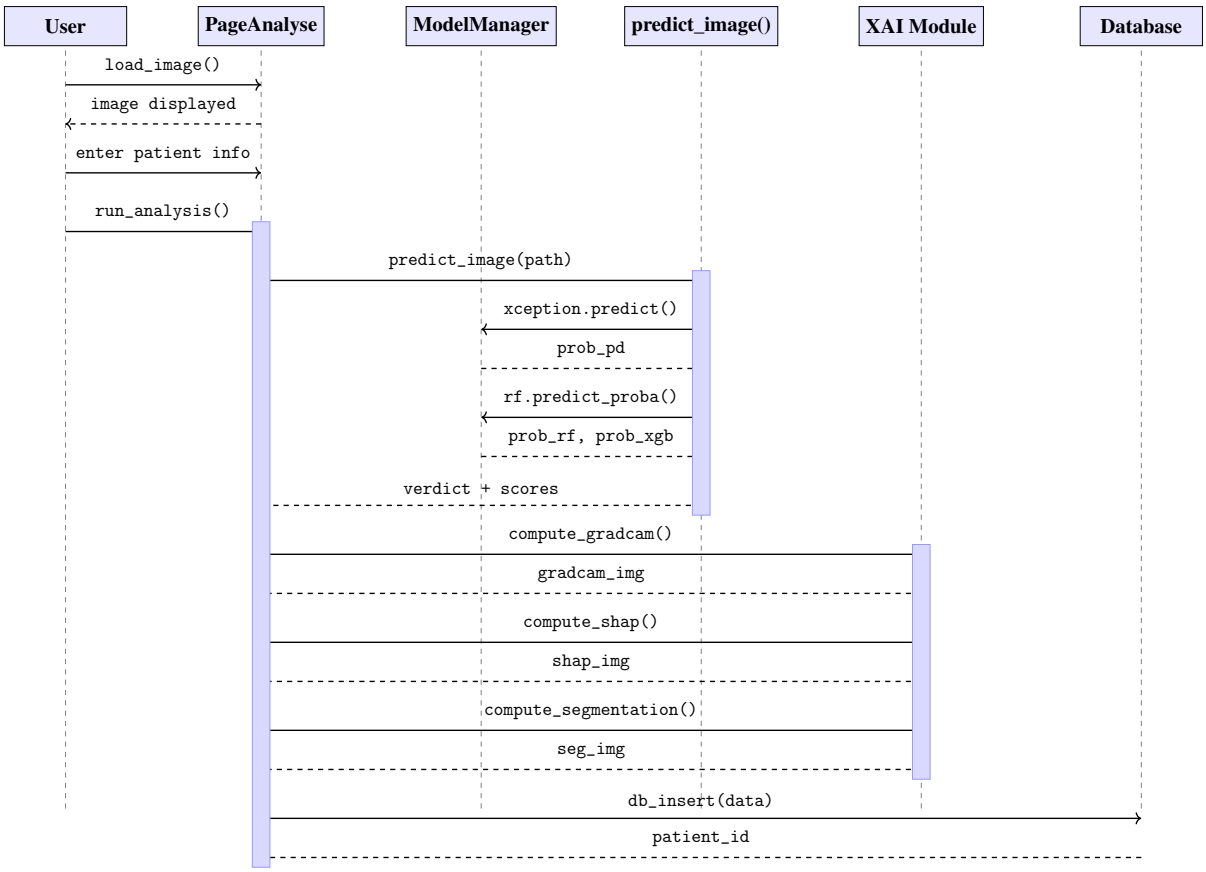


Figure 3.4: Sequence Diagram — Main Analysis Scenario

3.4 Dataset Presentation

3.4.1 Dataset Description

The NTUA Parkinson Dataset [23], which is openly accessible on the Github platform, served as the dataset used in this investigation. It is made up of about 43,000 brain MRI pictures that are separated into two groups: people with Parkinson's disease (PD) and people without the disease (healthy individuals). The pictures are arranged into two folders: PNG and JPEG versions of `PD_Patients` and `Non_PD_Patients` are accessible. This dataset was chosen because of its binary classification structure, which is consistent with the study's goal, and its comparatively large size, which is crucial for efficiently training deep learning models.

3.4.2 Class Distribution

There is a notable class imbalance in the sample, with a significantly higher proportion of non-PD photos than PD images. A class weighting method was used during training to overcome this problem, giving the Non-PD class a weight of 3.0 and the PD class a weight of 1.0 [25]. By preventing the model from becoming biased in favor of the majority class, this method enhances the model's capacity to accurately identify cases of Parkinson's disease. Out of the approximately 44,000 images in the dataset, 32,706 belong to the PD class and 10,381 belong to the Non-PD class, corresponding to a class ratio of approximately 3.15:1. This imbalance motivated the class weighting strategy described above.

Class	**Label**	**Class Weight**
Non-PD Patients	0	3.0
PD Patients	1	1.0
Total	—	43,087 images

Table 3.1: Dataset class distribution and applied weights

3.5 Data Preprocessing

3.5.1 Image Loading and Format Conversion

The `os` library was used to load the photos from the dataset directory by recursively searching the `PD_Patients` and `Non_PD_Patients` folders. Files with valid image extensions (.png,.jpg,.jpeg) were the only ones kept. To ensure reproducibility, the file paths and labels were placed in a Pandas DataFrame [26], which was then shuffled using a fixed random seed (`random_state=42`).

3.5.2 Image Resizing

To satisfy the Xception architecture's input requirements, all photos were downsized to a fixed resolution of 224×224 pixels [15]. The `ImageDataGenerator` class from the Keras library [27] carried out this resizing automatically.

3.5.3 Normalization

The built-in preprocessing function of Xception (`tf.keras.applications.xception.preprocess_inp`) which scales pixel values to the range $[-1, 1]$ [15], was used to normalize the pixel values. All

three subsets—training, validation, and test sets—were subjected to this normalization.

3.5.4 Data Augmentation

Data augmentation was only applied to the training set using the `ImageDataGenerator` class [27] in order to increase the model's capacity for generalization and decrease overfitting. Table 3.2 provides a summary of the augmentation parameters.

Augmentation Technique	**Parameter Value**
Rotation range	20°
Width shift range	0.2
Height shift range	0.2
Shear range	0.2
Zoom range	0.2
Horizontal flip	True
Fill mode	Nearest

Table 3.2: Data augmentation parameters applied to the training set

3.5.5 Dataset Splitting

To maintain the class distribution across all subsets, the dataset was split into three subgroups using a stratified split approach [25]. Table 3.3 displays the splitting ratio.

Subset	**Ratio**	**Approx. Size**
Training set	70%	$\approx$ 30,800 images
Validation set	15%	$\approx$ 6,600 images
Test set	15%	$\approx$ 6,600 images
Total	100%	$\approx$ 44,000 images

Table 3.3: Dataset splitting ratios

3.5.6 Class Imbalance Handling

The dataset has a class imbalance between PD and non-PD samples, as was noted in the preceding section. During model training, a class weighting method was used to solve this problem without changing the dataset itself. In particular, the model was penalized more severely for incorrectly classifying the minority class by giving the Non-PD class a weight of 3.0 and the PD class a weight of 1.0 [25].

3.6 Models Configuration

3.6.1 Experimental Environment

Every experiment was carried out on the Kaggle cloud platform, which offered two NVIDIA Tesla T4 GPUs with 16 GB of VRAM each for free. TensorFlow's `MirroredStrategy` [27] was used to manage the distributed training, synchronizing gradient changes across both GPUs to speed up the training process. Tables 3.4 and 3.5 provide a summary of the hardware specifications and software environment used in this investigation.

Component	**Specification**
Platform	Kaggle Cloud
GPU	2 × NVIDIA Tesla T4
VRAM	16 GB per GPU
RAM	30 GB
GPU Strategy	MirroredStrategy

Table 3.4: Hardware specifications

Library	**Version**
Python	3.10
TensorFlow	2.x
Keras	built-in (tf.keras)
Scikit-learn	1.x
XGBoost	1.x
SHAP	0.x
OpenCV	4.x
NumPy	1.x
Pandas	1.x
Matplotlib	3.x

Table 3.5: Software environment

3.6.2 Xception Configuration

The Xception model was selected as the main feature extractor and classifier in this study due to its effective use of depthwise separable convolutions, which reduce computational cost while maintaining high performance on image classification tasks. Given the comparatively large size of the NTUA Parkinson dataset, Strategy 2 (partial training) was adopted: the convolutional base, pre-trained on ImageNet, was kept frozen to preserve its general visual features, while only the custom classification head was retrained on the target dataset. Because of its demonstrated efficacy in picture classification tasks, the Xception architecture [15] was chosen

as the main deep learning model in this work. Using the transfer learning approach outlined in the previous chapter, the model was started with pre-trained weights from the ImageNet dataset. During the first training step, the learnt feature representations were preserved by freezing the base model. The original model was supplemented with a custom classification head that included a GlobalAveragePooling2D layer, two fully connected Dense layers with 256 and 128 neurons, respectively, each preceded by a Dropout layer with a rate of 0.3 to avoid overfitting. Two neurons with a Softmax activation function for binary classification make up the final output layer. With a learning rate of 0.001 and categorical crossentropy as the loss function, the model was assembled using the Adam optimizer [28]. Table 3.6 provides an overview of the classification head's architecture.

Layer	**Output Shape**	**Parameters**
Xception base (frozen)	$(7, 7, 2048)$	ImageNet weights
GlobalAveragePooling2D	$(2048,)$	—
Dense + ReLU	$(256,)$	trainable
Dropout	$(256,)$	rate = 0.3
Dense + ReLU	$(128,)$	trainable
Dropout	$(128,)$	rate = 0.3
Dense + Softmax	$(2,)$	trainable

Table 3.6: Xception classification head architecture

Three callbacks governed the training process: Early stopping while keeping an eye on the validation loss for five epochs, ModelCheckpoint, which saves the optimal model weights based on validation accuracy, and ReduceLROnPlateau, which lowers the learning rate by a factor of 0.5 when the validation loss stagnates for three consecutive epochs. Using the MirroredStrategy from TensorFlow to take advantage of the two available T4 GPUs on the Kaggle platform [27], the model was trained for a maximum of 30 epochs with a batch size of 64.

Hyperparameter	**Value**
Optimizer	Adam
Learning rate	0.001
Batch size	64
Maximum epochs	30
Early stopping	patience = 5
ReduceLROnPlateau	factor = 0.5, patience = 3
Loss function	Categorical Crossentropy
GPU strategy	MirroredStrategy (T4 $\times$ 2)

Table 3.7: Xception training hyperparameters

The trained Xception model was employed as a feature extractor in the second stage. In

particular, a feature vector with 2048 dimensions was created for every image by extracting the output of the GlobalAveragePooling2D layer. StandardScaler [25] was then used to standardize these features in order to guarantee a zero mean and unit variance. The feature dimensionality was then reduced from 2048 to 500 using the SelectKBest algorithm with the f_classif scoring function, keeping only the most statistically significant features [25].

3.6.3 Random Forest Configuration

Random Forest was selected as one of the traditional machine learning classifiers due to its robustness to overfitting and its ability to handle high-dimensional feature vectors extracted by Xception. The 500 chosen features were used to train the Random Forest classifier [16]. To further address the class imbalance, the model was set up with 500 decision trees, no maximum depth limitation, and the class_weight parameter set to `balanced`. For reproducibility, a fixed random seed of 42 was employed, and `n_jobs=-1` was used to make use of all CPU cores.

Hyperparameter	**Value**
Number of estimators	500
Maximum depth	None (unlimited)
Class weight	balanced
Random state	42
n_jobs	-1 (all cores)

Table 3.8: Random Forest hyperparameters

3.6.4 XGBoost Configuration

XGBoost was chosen as a second complementary classifier for its efficiency and strong performance on structured feature vectors, in order to compare its behavior with Random Forest under the same feature set. The 500 chosen features were also used to train the XGBoost classifier [17]. The model was set up using 300 estimators, a learning rate of 0.05, and a maximum tree depth of 5. To lessen overfitting, the subsampling settings for both `subsample` and `colsample_bytree` were set to 0.8. To address class imbalance, the `scale_pos_weight` parameter was set to the ratio of the number of PD and non-PD samples. For computational efficiency, the histogram-based tree technique (`hist`) was employed.

Hyperparameter	Value
Number of estimators	300
Maximum depth	5
Learning rate	0.05
Subsample	0.8
Colsample bytree	0.8
Min child weight	3
Gamma	0.1
Scale pos weight	auto (class ratio)
Tree method	hist
Random state	42

Table 3.9: XGBoost hyperparameters

3.6.5 Final Decision — Majority Voting

A majority vote system was used to merge the predictions of the three models (Xception, Random Forest, and XGBoost) into a single, more reliable diagnosis. A unique decision threshold was empirically established for each classifier based on its output probability distribution because each model was trained and calibrated independently: 0.50 for Xception, 0.70 for Random Forest, and 0.85 for XGBoost. Each model generates a likelihood score for the PD class for a specific patient. A model is deemed to "vote" for the PD class if its predicted probability reaches or exceeds its associated threshold; otherwise, it votes for the Non-PD class. The patient is classed as "Parkinson Detected" if at least two of the three models vote for PD; if not, they are labeled as "Normal." The final diagnosis is decided by majority voting. As evidenced by the confidence level shown in the ParkinsonAI application (see Section 4.6.1), this method lowers the possibility that a single model's incorrect classification would influence the final choice and raises the overall reliability of the diagnosis.

3.7 Conclusion

The entire process used to identify Parkinson's disease from brain MRI pictures was described in this chapter. The NTUA Parkinson dataset, which comprises roughly 44,000 MRI pictures split into two classes—PD patients and non-PD patients—was the first thing we introduced. The preprocessing pipeline, which included image resizing, normalization, data augmentation, stratified splitting, and class imbalance handling using class weighting, was then thoroughly explained. The configuration of every model used in this study was then explained

in detail: the Xception CNN served as the main classifier, and after dimensionality reduction using SelectKBest, it was used as a feature extractor feeding two traditional machine learning models, Random Forest and XGBoost. Finally, the dual T4 GPU experimental environment built on the Kaggle platform was showcased. The next chapter will go into great detail on the outcomes of these models as well as the explainability analysis using Grad-CAM, SHAP, and segmentation techniques.

Chapter 4

Experimentals Results

4.1 Introduction

The experimental findings from the suggested Parkinson's disease detection method are presented and discussed in this chapter. The Xception convolutional neural network is combined with two classifiers, Random Forest classifier, the XGBoost classifier. Standard evaluation criteria, including as accuracy, precision, recall, and F1-score, are used to evaluate each model's performance. Additionally, the outcomes of the explainable AI methods, Grad-CAM, SHAP and segmentation, are shown and examined. Lastly, the created desktop application is presented together with a demonstration of its primary features.

4.2 Evaluation Metrics

To evaluate the effectiveness of the models employed in this research, a number of common assessment indicators were used. These measurements are calculated using the confusion matrix, a summary of the categorization findings into four groups: True Negatives (TN), True Positives (TP), False Negatives (FN) and False Positives (FP).

Accuracy measures the overall proportion of correctly classified samples:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \tag{4.1}$$

Precision measures the proportion of correctly predicted positive cases among all predicted positive cases:

$$Precision = \frac{TP}{TP + FP} \tag{4.2}$$

Recall (Sensitivity) measures the proportion of actual positive cases that were correctly identified:

$$Recall = \frac{TP}{TP + FN} \tag{4.3}$$

F1-Score is the harmonic mean of precision and recall, providing a balanced measure especially useful in the presence of class imbalance:

$$F1\text{-}Score = 2 \times \frac{Precision \times Recall}{Precision + Recall} \tag{4.4}$$

AUC-ROC (Area Under the Receiver Operating Characteristic Curve) measures the ability of the model to distinguish between classes across all classification thresholds [29]. These metrics are particularly relevant in the medical context of this study, where minimizing false negatives (missed Parkinson's cases) is of critical importance.

4.3 Xception Results

A test set of 6,602 MRI pictures was used to assess the performance of Xception. With a loss of 0.3169, the model's total accuracy was **84.19%**. Table 4.1 displays the obtained results in terms of precision, recall and F1-score.

Class	Precision	Recall	F1-Score	Support
Non-PD (Sain)	0.64	0.82	0.72	1607
PD (Parkinson)	0.94	0.85	0.89	4995
Accuracy		0.84		6602
Macro avg	0.79	0.83	0.80	6602
Weighted avg	0.86	0.84	0.85	6602

Table 4.1: Xception classification report on the test set

The results show that the Xception model has a high precision of 0.94 for the PD class,

which means that 94% of the time, the model correctly predicts Parkinson's disease. The model correctly detects 85% of real Parkinson's patients, according to the PD class recall of 0.85. The substantial class imbalance in the dataset, where PD samples (32,706) significantly outnumber Non-PD samples (10,381), is the reason for the comparatively lower precision for the Non-PD class (0.64).

The confusion matrix derived from the Xception model on the test set is shown in Figure 4.1.

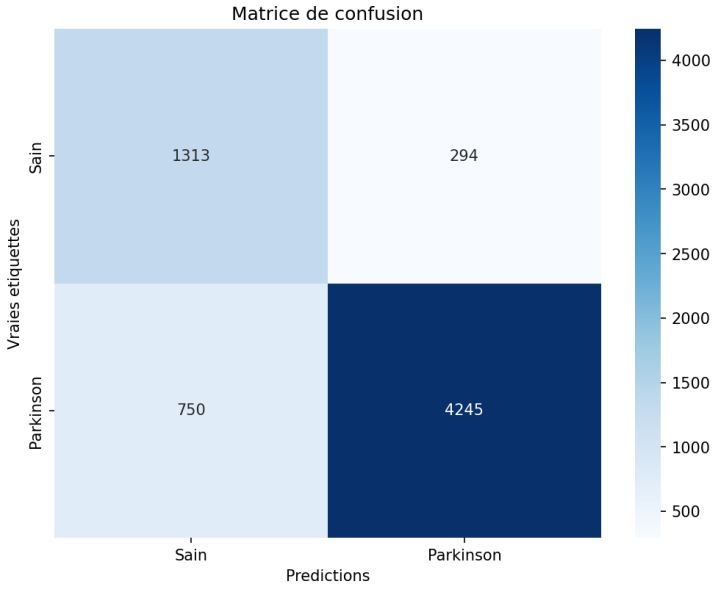


Figure 4.1: Confusion matrix of the Xception model

Figure 4.2 presents the ROC curves of all three models, illustrating their ability to discriminate between PD and Non-PD classes across different classification thresholds [29].

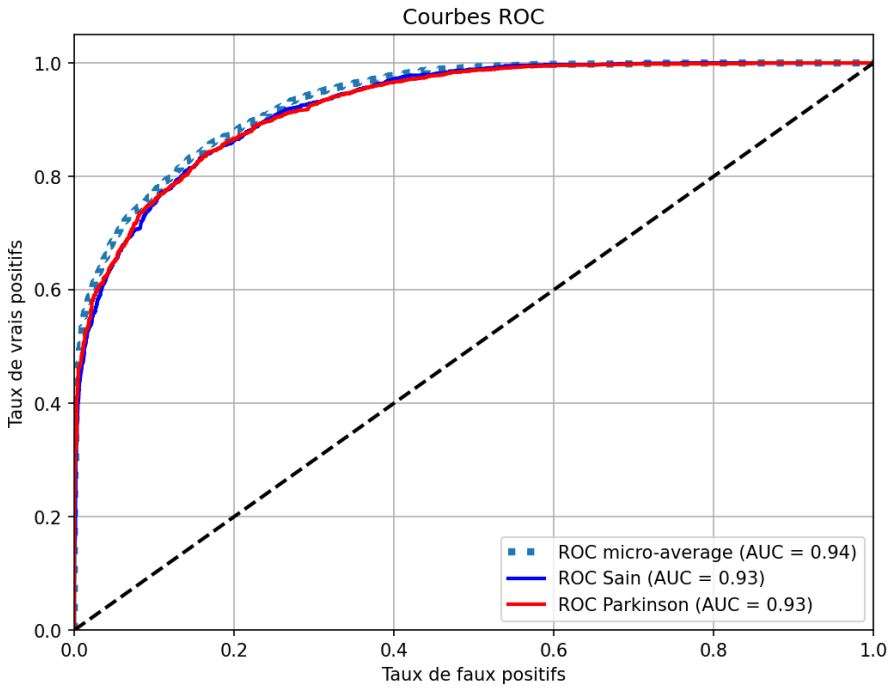


Figure 4.2: ROC curves of Xception, Random Forest and XGBoost

4.4 Random Forest and XGBoost Results

After the feature extraction stage, where the GlobalAveragePooling layer of the Xception model generated 2048-dimensional feature vectors that were then reduced to 500 using SelectKBest, two traditional machine learning models were trained and evaluated on the test set.

4.4.1 Random Forest

On the test set, the Random Forest classifier achieved an accuracy of **81.40%**. The comprehensive classification report is shown in Table 4.2.

Class	Precision	Recall	F1-Score	Support
Non-PD (Sain)	0.98	0.24	0.39	1607
PD (Parkinson)	0.80	1.00	0.89	4995
Accuracy		0.81		6602
Macro avg	0.89	0.62	0.64	6602
Weighted avg	0.85	0.81	0.77	6602

Table 4.2: Random Forest classification report on the test set

With a perfect recall of 1.00 for the PD class and a very high precision of 0.98 for the Non-PD class, the Random Forest model is able to identify every instance of Parkinson's disease. However, the poor recall of 0.24 for the Non-PD class shows that a large number of healthy individuals are incorrectly categorized as PD patients, which illustrates how the class imbalance affects the behavior of the model [16].

4.4.2 XGBoost

On the test set, the XGBoost classifier achieved an accuracy of **79.13%**. The comprehensive classification report is shown in Table 4.3.

Class	**Precision**	**Recall**	**F1-Score**	**Support**
Non-PD (Sain)	0.99	0.14	0.25	1607
PD (Parkinson)	0.78	1.00	0.88	4995
Accuracy		0.79		6602
Macro avg	0.89	0.57	0.56	6602
Weighted avg	0.83	0.79	0.73	6602

Table 4.3: XGBoost classification report on the test set

Like Random Forest, XGBoost attains a perfect recall of 1.00 for the PD class and a near-perfect precision of 0.99 for the Non-PD class. Nevertheless, the extremely poor recall of 0.14 for the Non-PD class indicates that the threshold value of 0.85 used during inference has a direct impact on the model's tendency to classify the majority of samples as PD [17].

4.4.3 Models Comparison

The accuracy of the three models is visually compared in Figure 4.3, which gives a clear picture of how well each model performs in relation to the others.

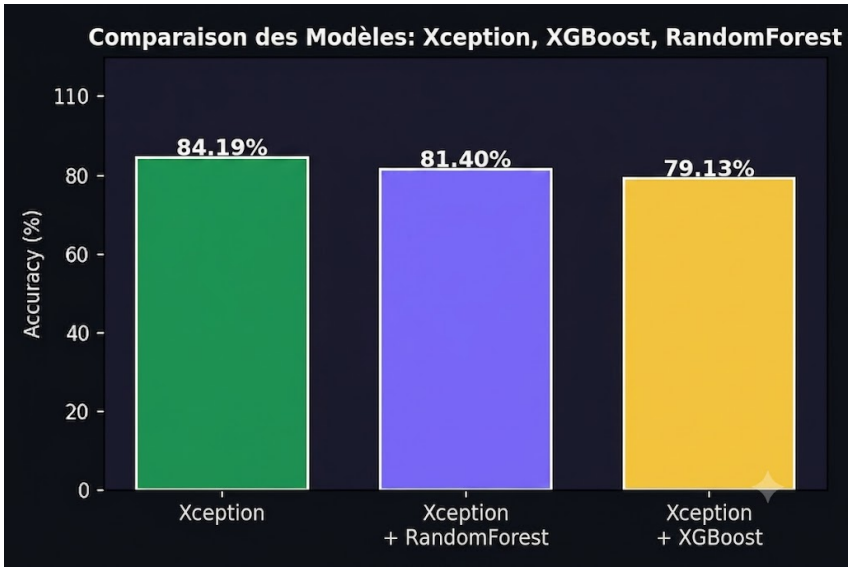


Figure 4.3: Accuracy comparison of Xception, Random Forest and XGBoost

Overall, the Xception CNN achieves the highest accuracy of 84.19% while maintaining a more balanced precision-recall trade-off across both classes, outperforming both traditional machine learning models.

4.5 XAI Results

4.5.1 Grad-CAM Results

The Xception model was used to create visual explanations of its predictions using Gradient-weighted Class Activation Mapping (Grad-CAM) [19]. The heatmaps show the areas of the brain that most affected the model's choice, utilizing a color scale that went from red (high activation) to blue (low activation). A comparison of Grad-CAM heatmaps for PD and Non-PD MRI images is shown in Figure 4.4.

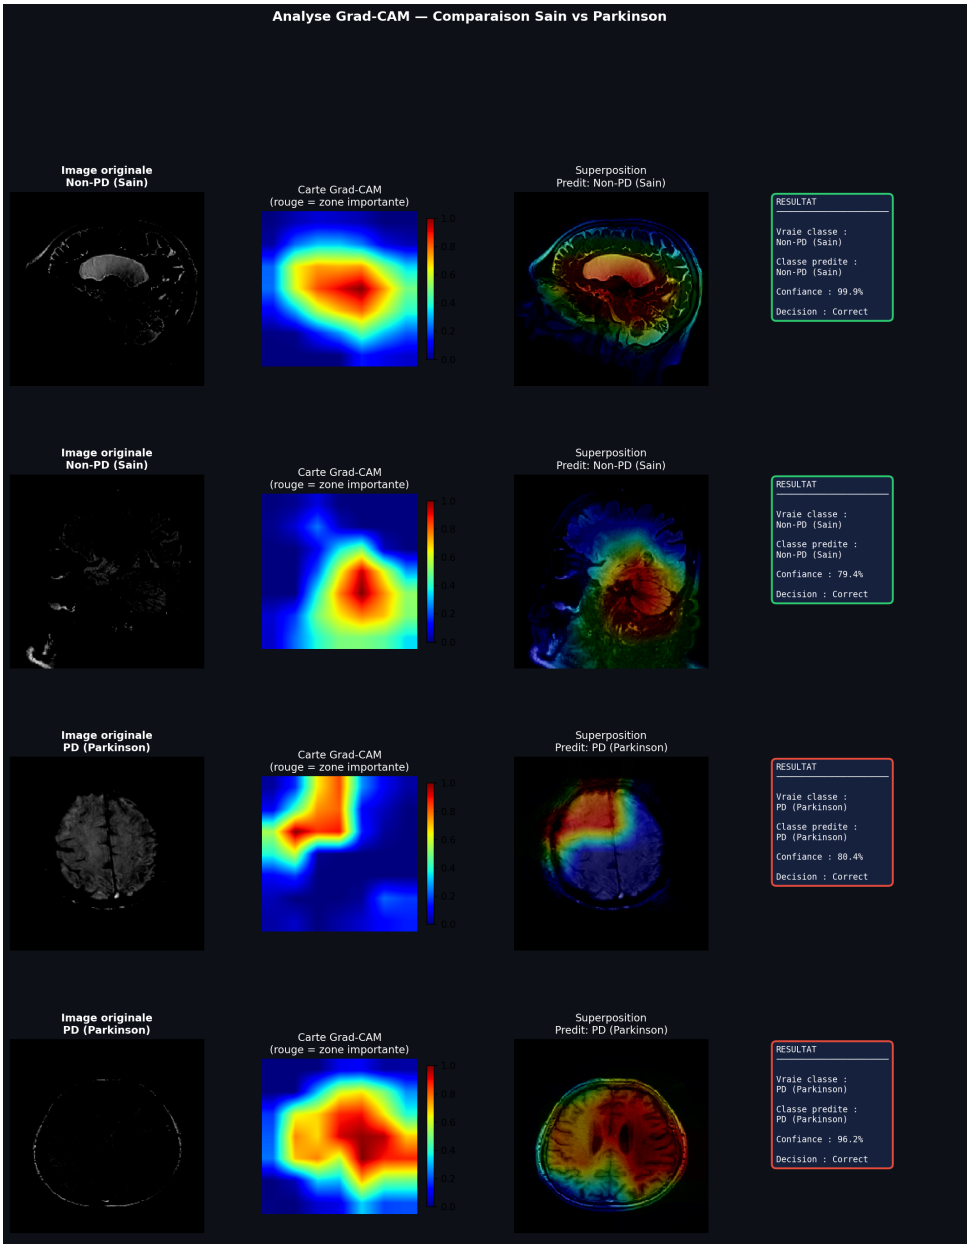


Figure 4.4: Grad-CAM comparative heatmaps for PD and Non-PD cases

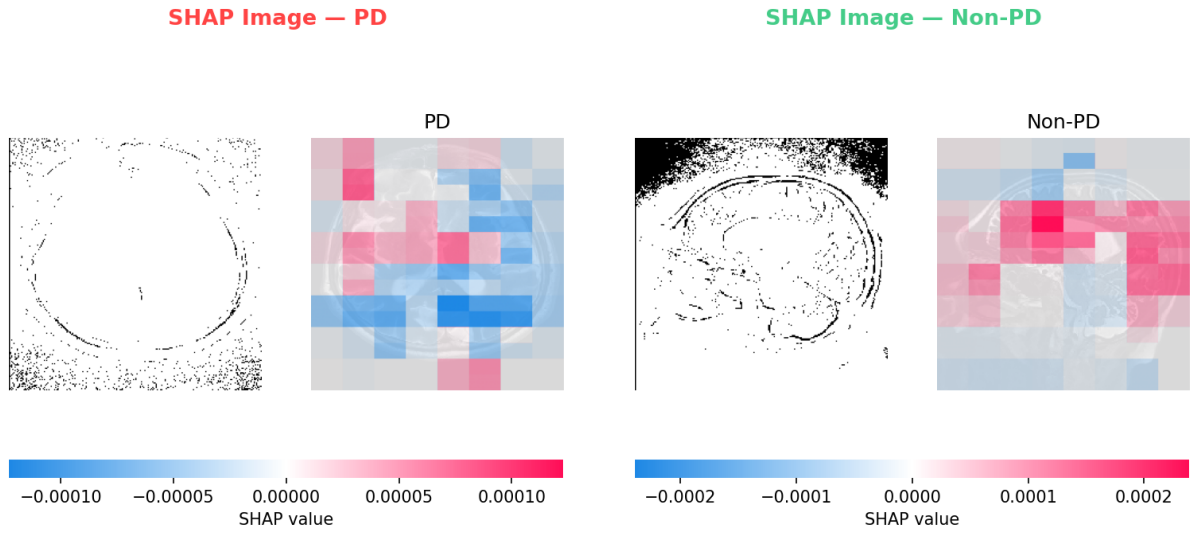
The red zones for PD patients represent aberrant regions found by the model, especially in areas connected to the Basal Ganglia and Substantia Nigra, which are known to be impacted in Parkinson's disease [2]. The activations are much weaker in Non-PD individuals, suggesting that there are no structural abnormalities in the brain. These findings verify that the model bases its predictions on areas that are clinically significant.

4.5.2 SHAP Results

To guarantee a thorough interpretation of the model's choices, SHAP (SHapley Additive exPlanations) [21] is used in this study in two complementary methods.

4.5.2.1 SHAP Image Explainer on MRI

Using a blur masker to examine the contribution of various image areas to the model's prediction, the first method applies SHAP directly to the MRI images processed by the Xception CNN. Red and pink-highlighted areas in each image pushed the model in the direction of the anticipated class, whereas blue areas pushed in the other direction.



(a) SHAP Image — PD case

(b) SHAP Image — Non-PD case

Figure 4.5: SHAP Image Explainer — PD vs Non-PD brain MRI regions

The red/pink areas for Parkinson's disease (PD) instances represent brain regions that significantly affected the model's choice to identify Parkinson's disease, especially in areas connected to the Basal Ganglia and Substantia Nigra. The highlighted areas for non-PD instances represent visual characteristics that are not present in PD patients.

4.5.2.2 SHAP on Clinical Data

In the second method, a Random Forest classifier trained on clinical data taken from the NTUA Parkinson dataset CSV files is subjected to a `TreeExplainer` [21]. Age, UPDRS I, UPDRS II, UPDRS III, MMSE, FOG, and sex are the seven clinical characteristics that are examined.

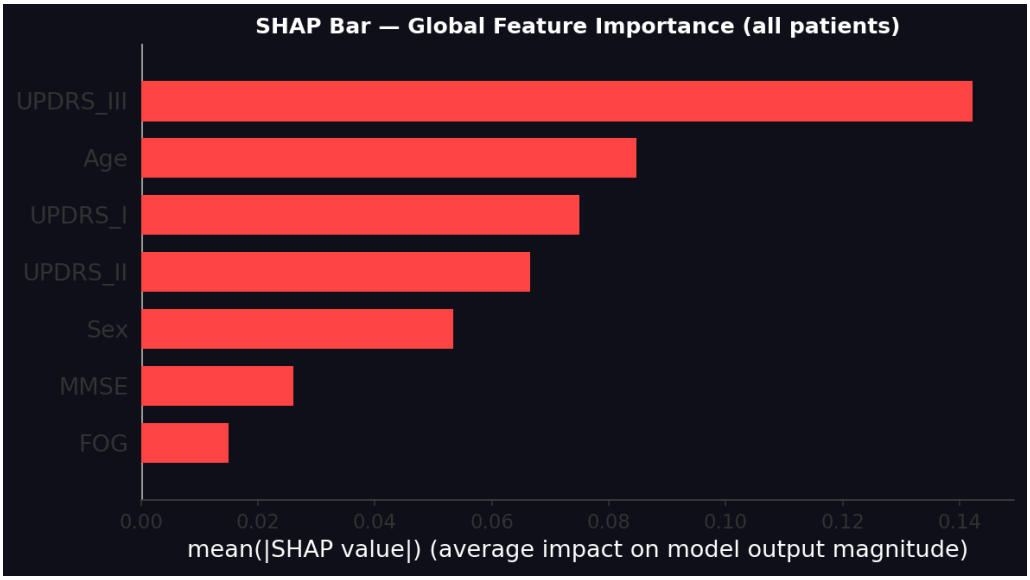


Figure 4.6: SHAP Bar plot — global clinical feature importance across all patients

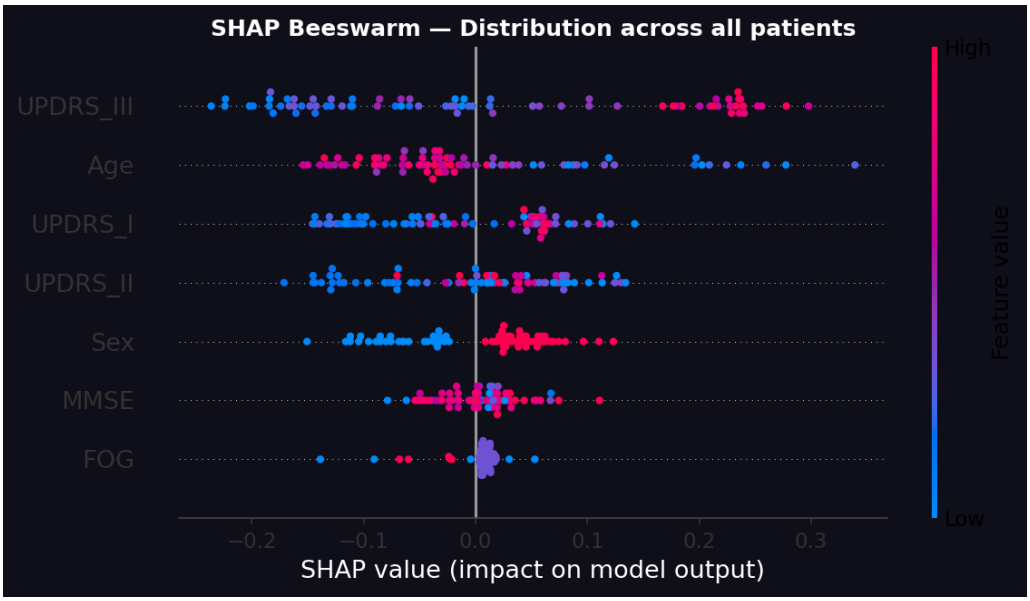


Figure 4.7: SHAP Beeswarm plot — SHAP value distribution across all patients

The Beeswarm figure (Figure 4.7) displays the distribution of SHAP values for all patients, with red dots denoting high feature values and blue dots denoting low feature values. The prediction is pushed toward the PD class by features with positive SHAP values and the Non-PD class by those with negative values. The Bar plot (Figure 4.6), which ranks the characteristics based on their mean absolute SHAP value across all patients, offers a global evaluation of clinical feature significance.

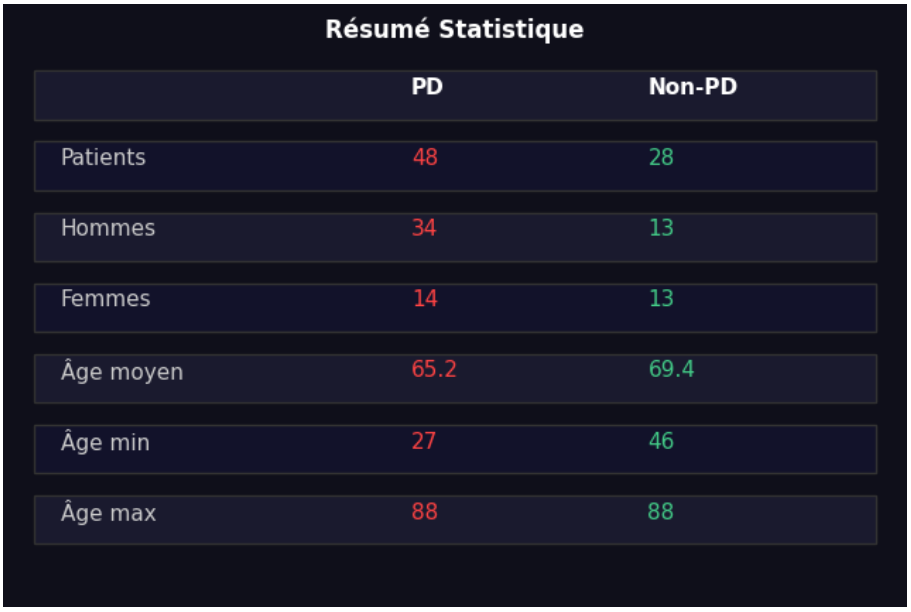


Figure 4.8: Statistical summary — comparison between PD and Non-PD patients

The statistical summary (Figure 4.8) confirms the clinical significance of the characteristics employed in the SHAP analysis by comparing important clinical indicators between PD and non-PD patients, such as age distribution and sex ratio.

4.5.3 Segmentation Results

The brain areas most linked to Parkinson's disease were identified and isolated using a segmentation technique based on Grad-CAM heatmaps [19]. The segmentation pipeline consists of four stages: the original MRI image, the Grad-CAM heatmap, the binary mask derived from an adaptive threshold, and the final segmentation overlay.

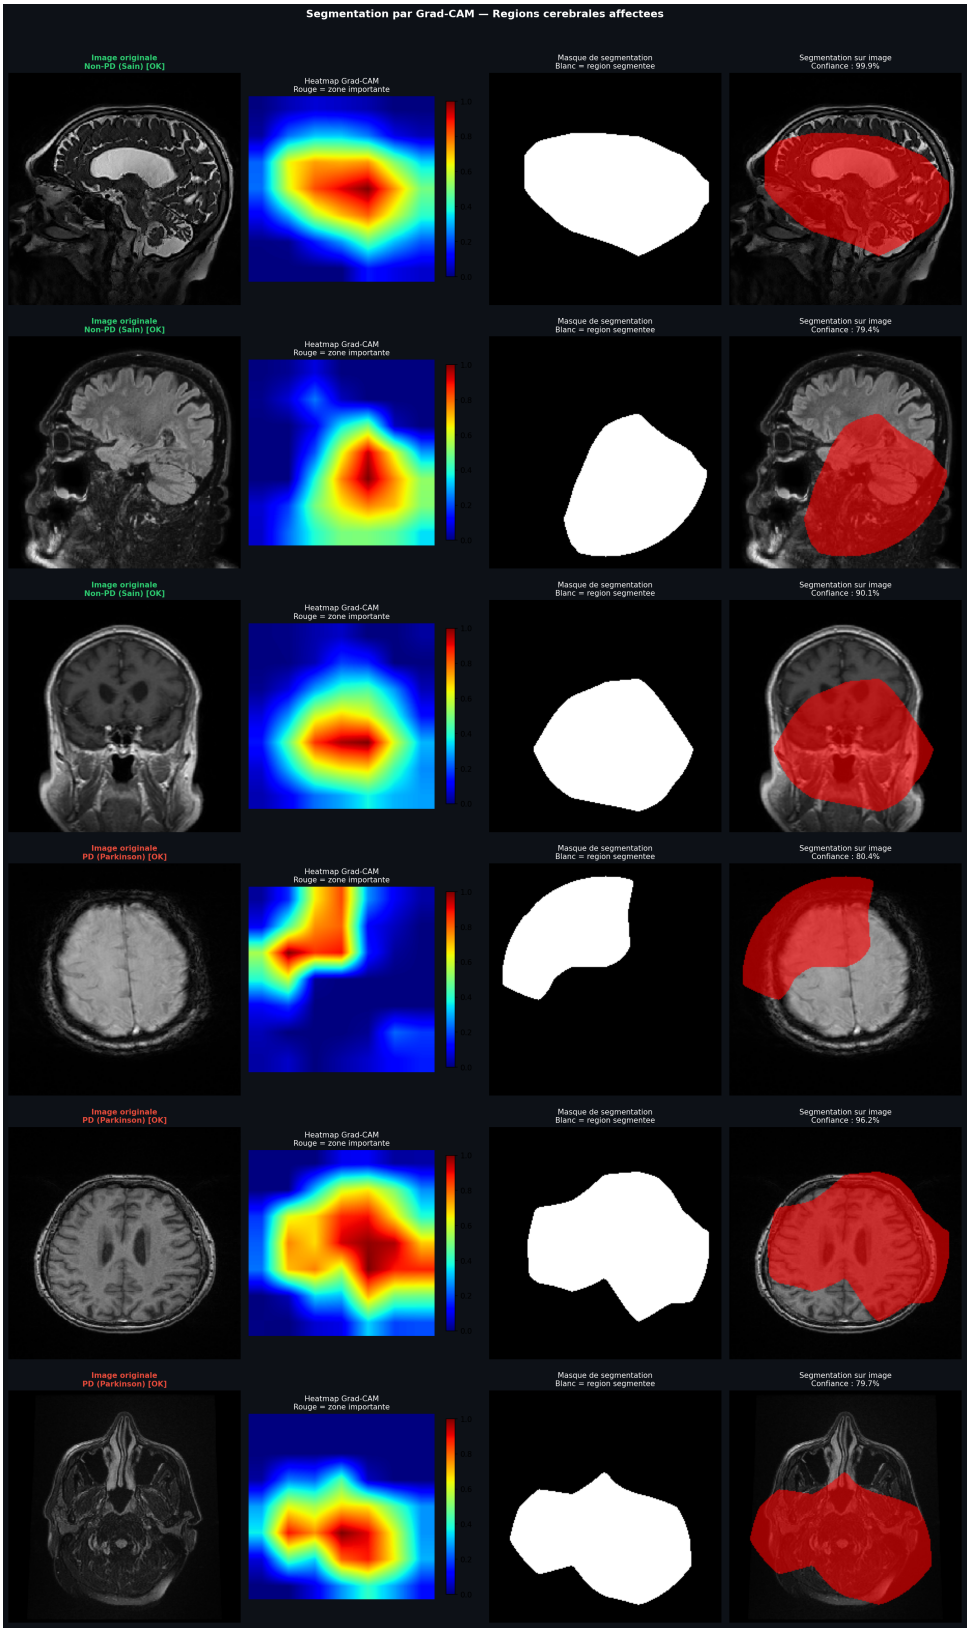


Figure 4.9: Segmentation results based on Grad-CAM heatmaps

The red-highlighted segmented regions for PD patients correlate to aberrant brain regions identified by the model, especially in the Substantia Nigra and Basal Ganglia regions [2]. The

segmented sections are small in Non-PD patients, indicating that there are no notable abnormalities. The consistency of the segmentation technique is demonstrated by the mean segmented area, which comprises around 27.57% of the brain region for PD cases and 27.95% for Non-PD cases.

4.6 Application Presentation

The CustomTkinter framework with a dark blue style was used to create the ParkinsonAI desktop application, which integrated the whole detection process into an easy-to-use interface. The three primary pages of the application are Analysis, History, and About.

4.6.1 Analysis Page

The application's primary user interface is represented by the Analysis page, which enables the user to import a brain MRI picture, input patient data, and execute the whole detection pipeline. The interface is separated into two panels, as seen in Figure 4.10. The verdict, model scores, confidence level, SHAP explanation, and brain segmentation results are displayed in the right panel, while the patient information form, the MRI image loader with dual preview (Original MRI and Grad-CAM), and the Run Analysis button are located in the left panel.

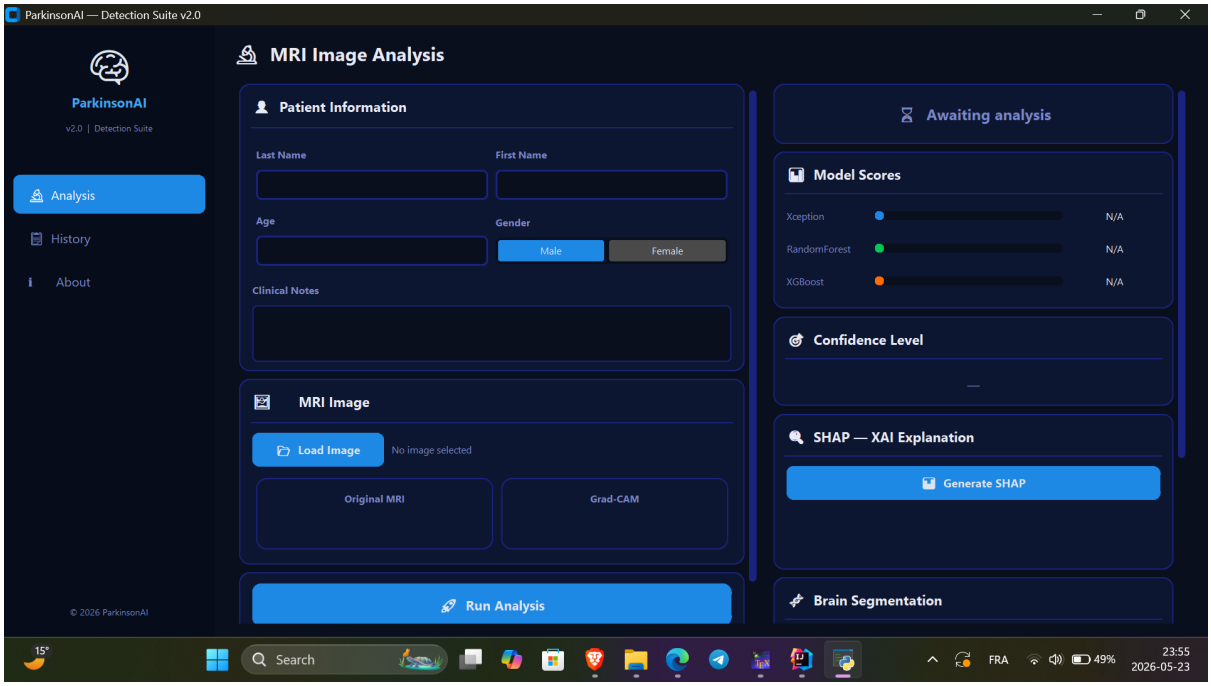


Figure 4.10: ParkinsonAI — Main Analysis page (awaiting analysis)

The application is shown in Figure 4.11 following the analysis of a PD MRI picture. "Parkinson Detected" is shown in red on the verdict banner, and the three model scores—Xception (66.5%), RandomForest (88.6%), and XGBoost (95.9%)—are displayed as progress bars. The diagnosis is confirmed with a 100% confidence level by the weighted vote score of 3/3.

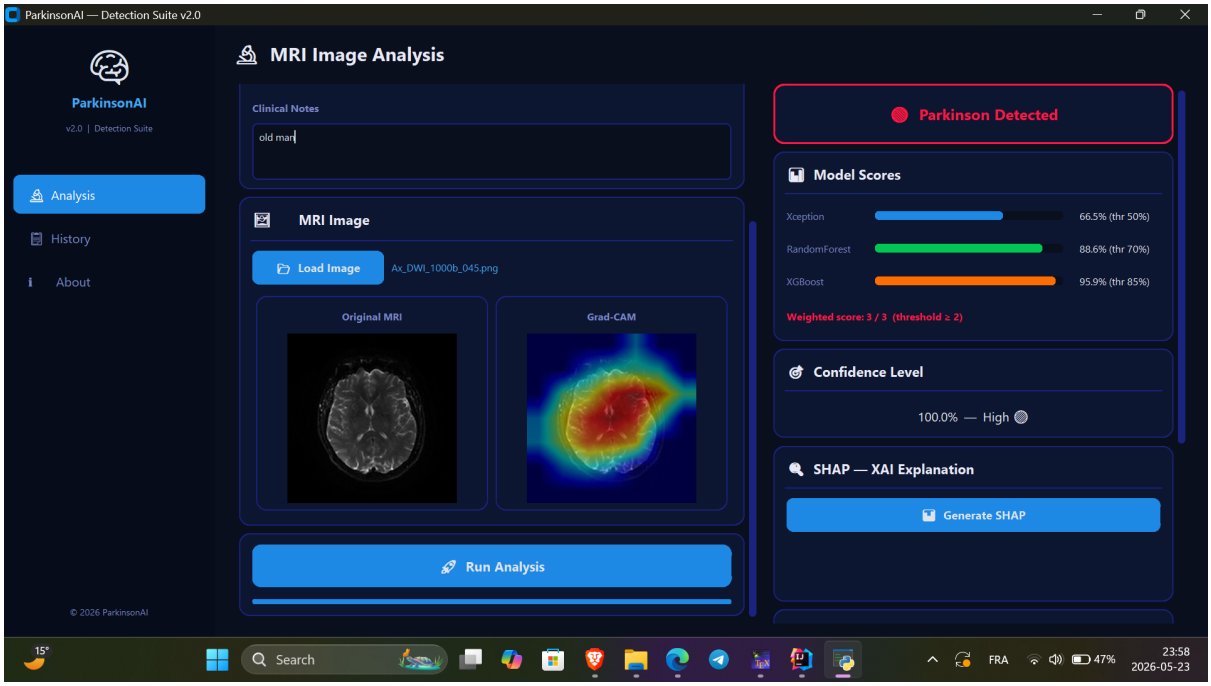


Figure 4.11: ParkinsonAI — Analysis result showing Parkinson Detected

The application's XAI findings, including the SHAP Waterfall plot and the Brain Segmentation with four panels (original picture, Grad-CAM heatmap, binary mask, and segmentation overlay), are shown in Figure 4.12.

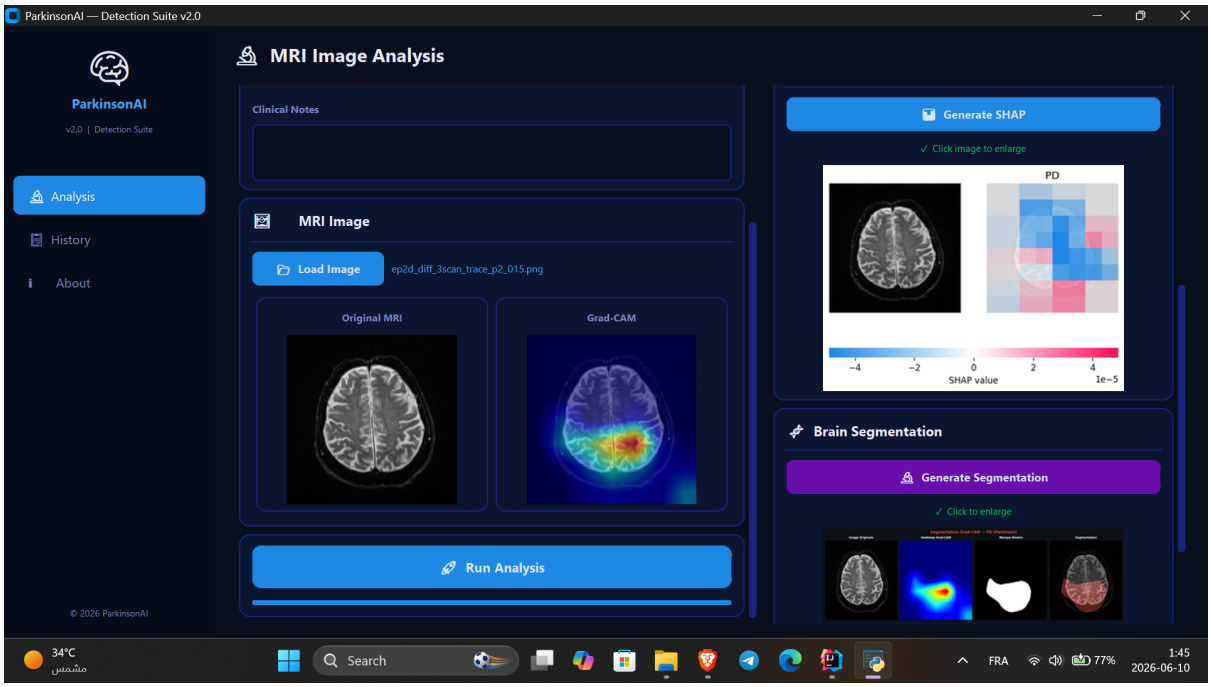


Figure 4.12: ParkinsonAI — SHAP and Segmentation results

4.6.2 History Page

As seen in Figure 4.13, the History tab offers a comprehensive list of all recorded patient records. The patient ID, first and last names, age, examination date, and verdict are all shown in each record. The user may see any patient's detailed view, delete specific data, and refresh the list.

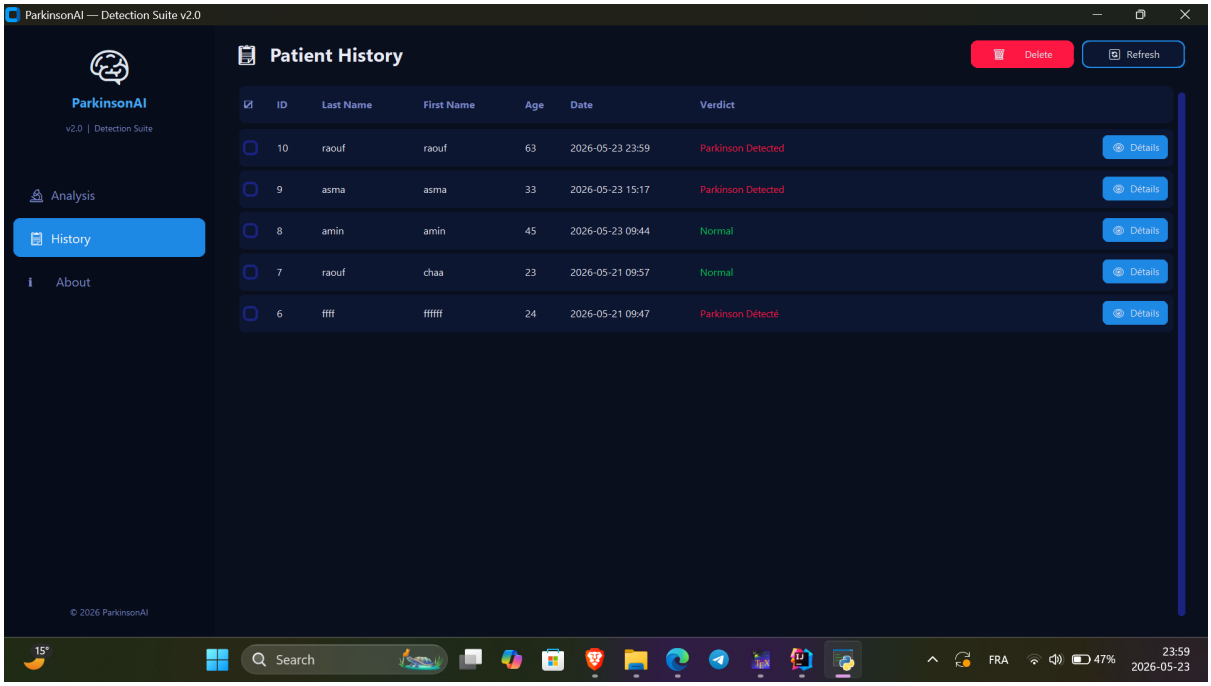


Figure 4.13: ParkinsonAI — Patient History page

The patient information, model scores, MRI and Grad-CAM pictures, segmentation results, SHAP explanation, and the ability to export a PDF report or modify the patient record are all included in the detailed patient view shown in Figure 4.14.



Figure 4.14: ParkinsonAI — Patient detail view with model scores

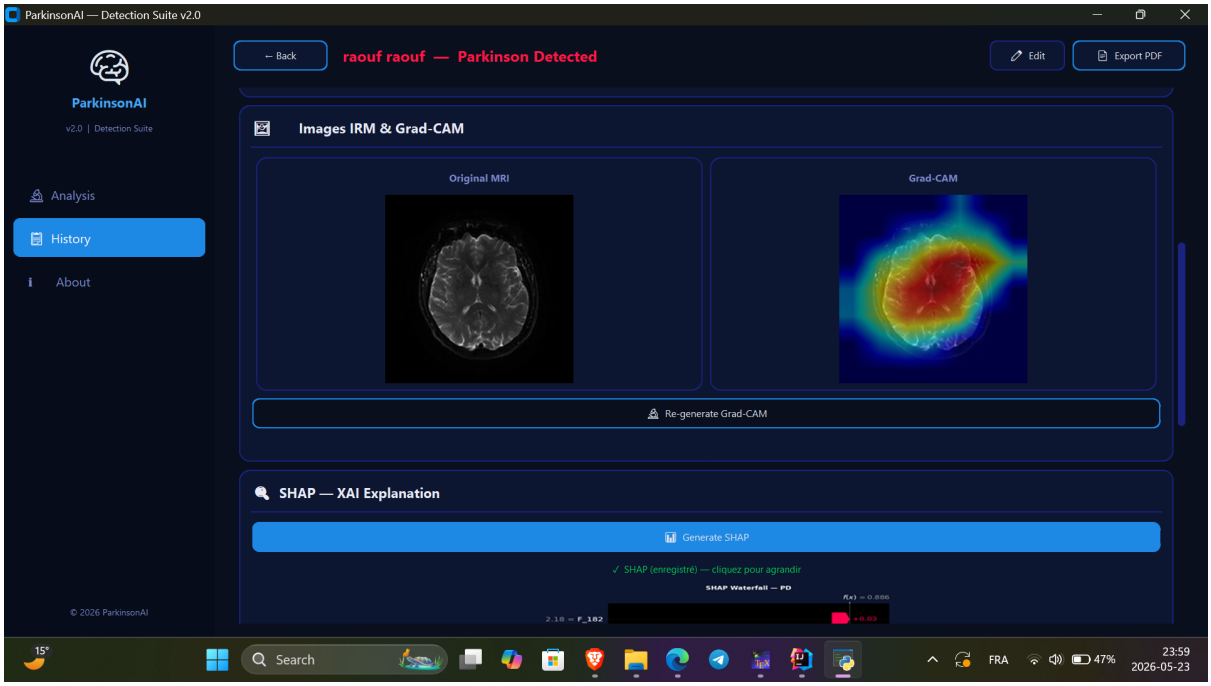


Figure 4.15: ParkinsonAI — Patient detail view with MRI and Grad-CAM

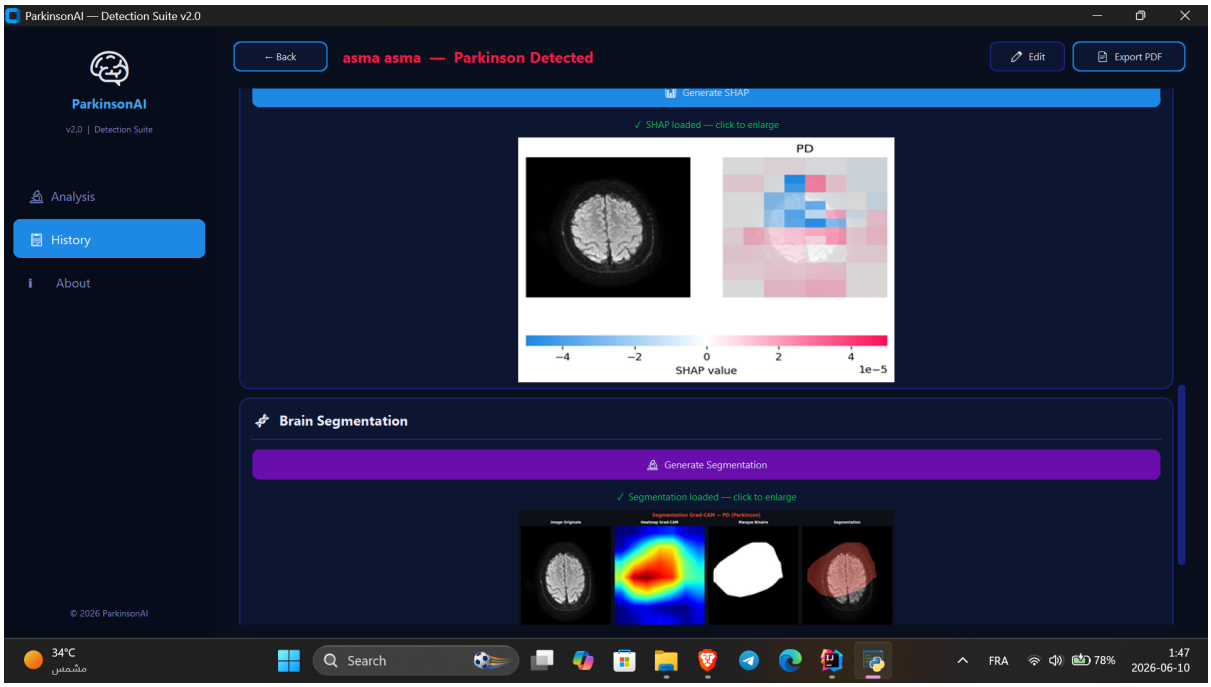


Figure 4.16: ParkinsonAI — Patient detail view with SHAP and Segmentation

4.6.3 Medical Report

As seen in Figure 4.17, the program offers an integrated medical report generator. The report contains clinical comments, model scores with thresholds, XAI analysis availability,

diagnosis verdict, and patient data. Using a web browser, the report may be printed straight from the program.

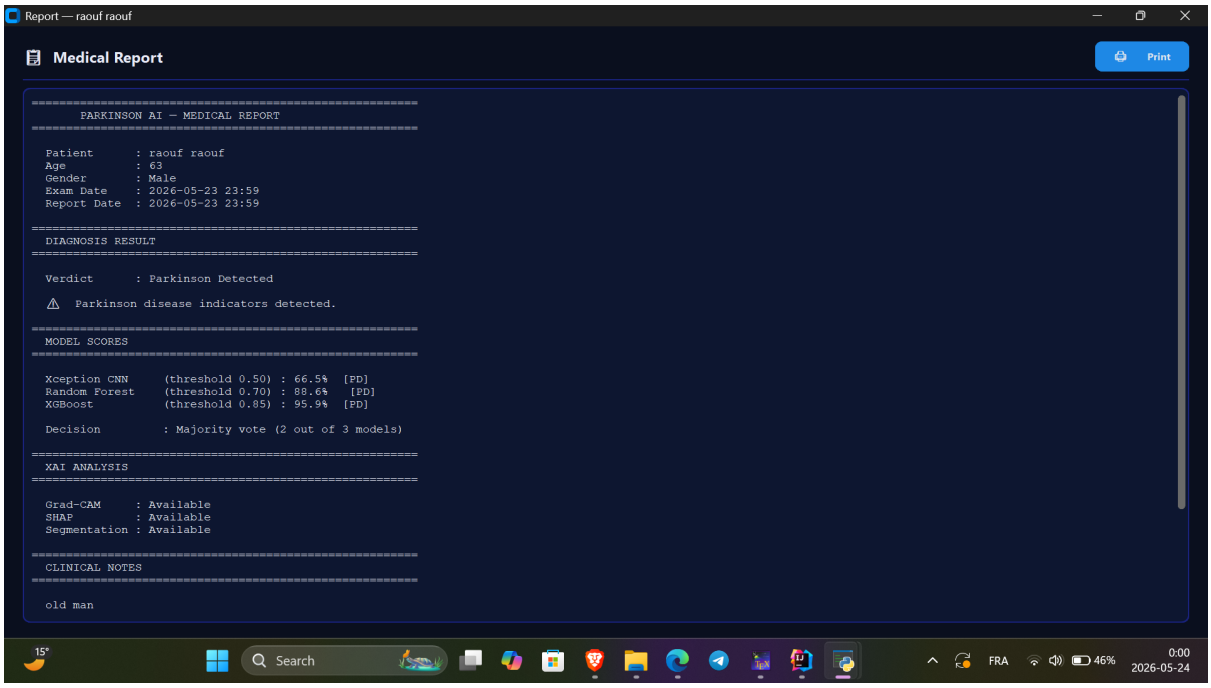


Figure 4.17: ParkinsonAI — Medical Report viewer

4.6.4 Edit Patient

As seen in Figure 4.18, the application enables the user to modify patient data at any moment via a specific popup window. The last name, first name, age, gender, and clinical notes may all be changed by the user, who can then save the modifications straight to the SQLite database.

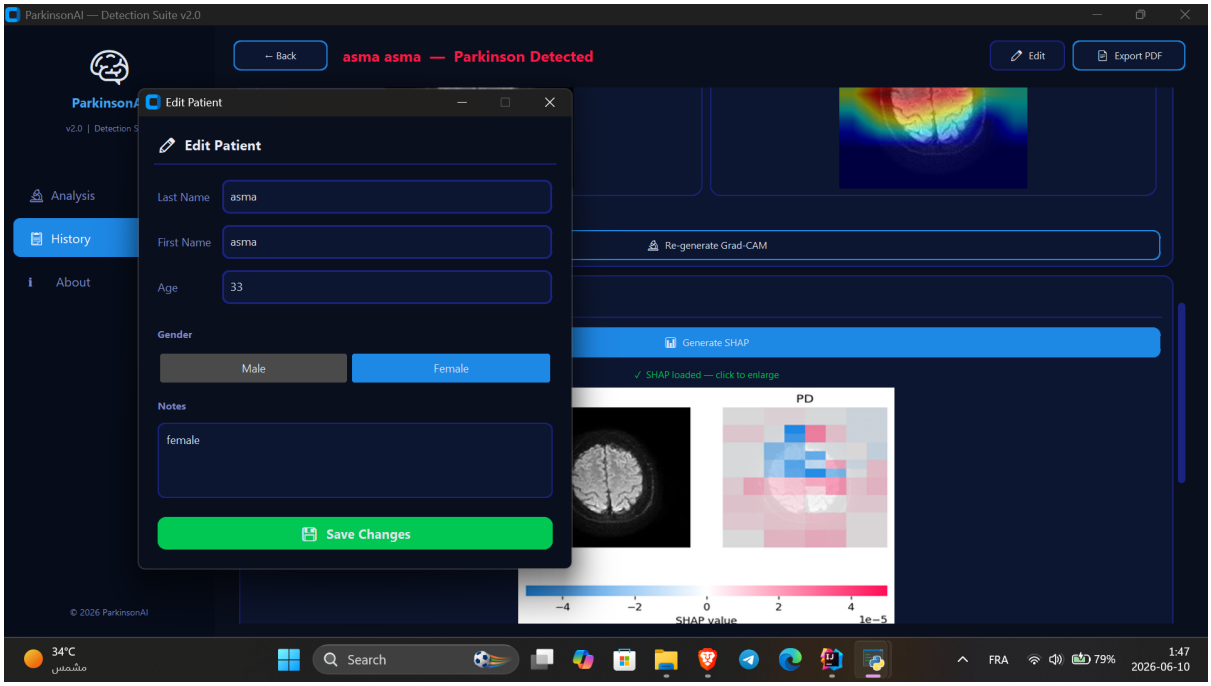


Figure 4.18: ParkinsonAI — Edit Patient popup

4.6.5 About Page

As seen in Figure 4.19, the About page provides a comprehensive overview of the program, including the AI models utilized, the Grad-CAM XAI approach, the technologies used, and an essential disclaimer addressing the tool's research-only function.

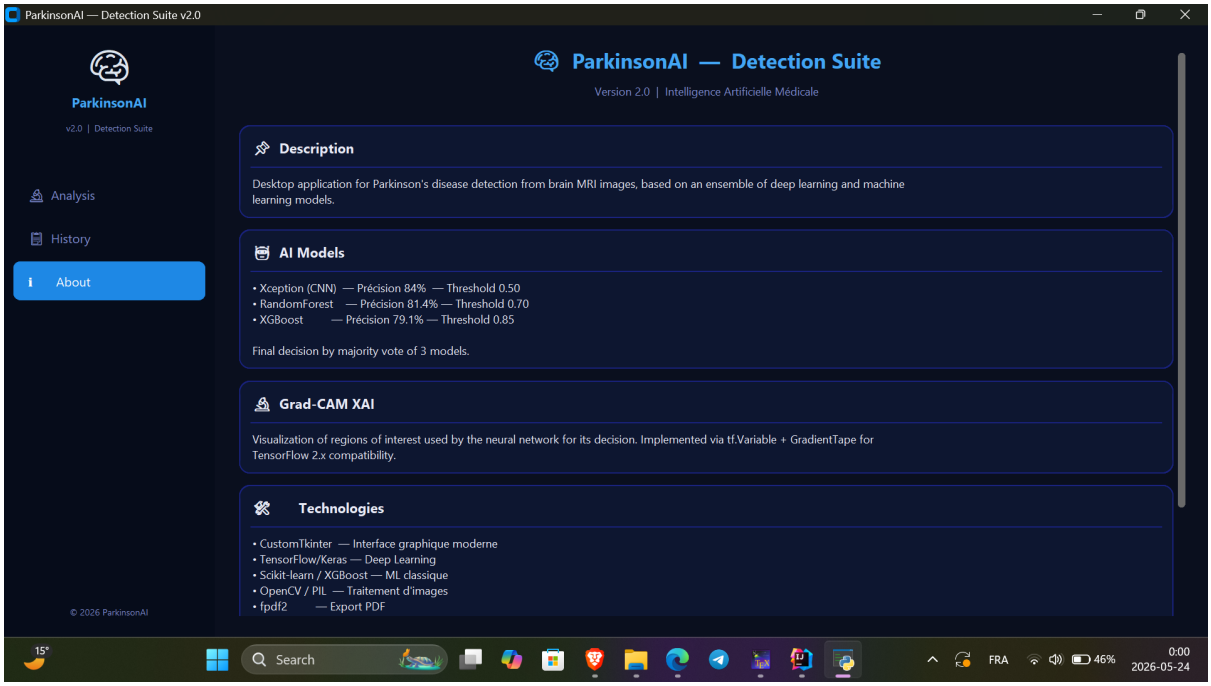


Figure 4.19: ParkinsonAI — About page

4.7 Commercial Web Platform

To showcase and distribute the ParkinsonAI software to researchers and medical professionals, a commercial online platform was created utilizing PHP and JavaScript in addition to the desktop version. The platform is organized into many sections that cover the whole commercial procedure, from license distribution to product presentation. Here is a link: https://parkinsonai.site.je

The platform's home page, seen in Figure 4.20, describes the program as an AI Medical Detection Platform for Parkinson's disease detection utilizing Xception CNN and XAI algorithms.

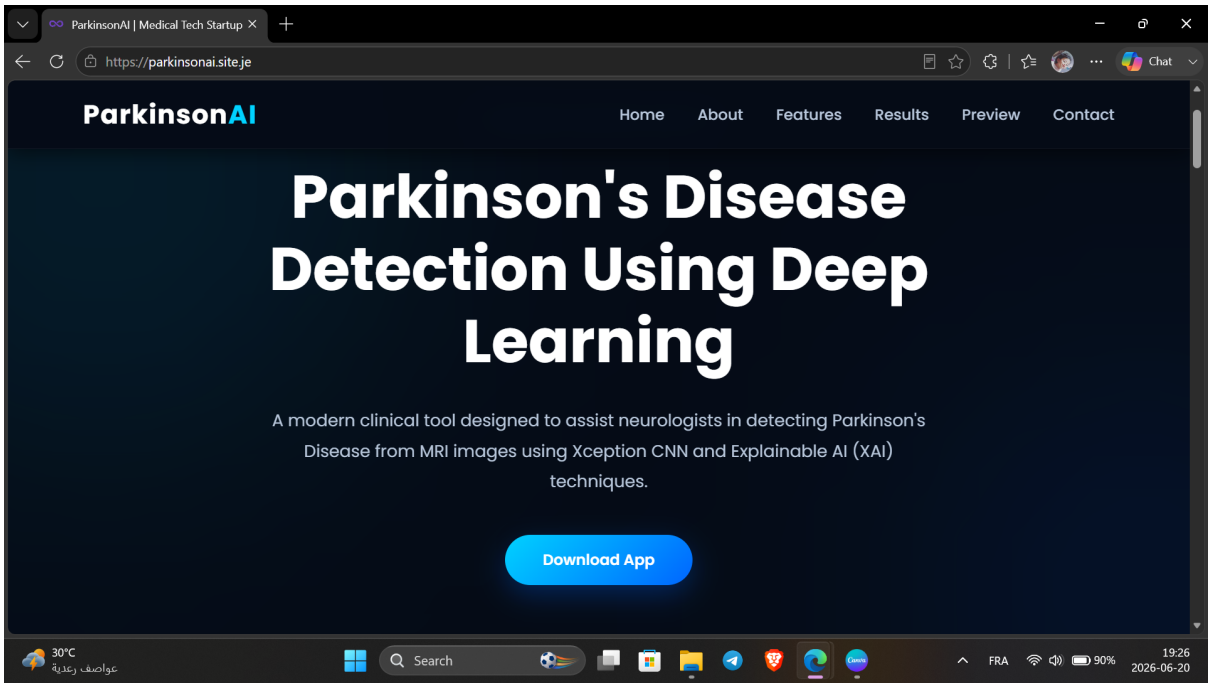


Figure 4.20: ParkinsonAI Web Platform — Home page

As seen in Figure 4.21, the platform provides two subscription plans and the Medical Clinic Plan, which targets hospitals with multi-user access and PACS/EHR system integration, and the Researcher Plan, which targets academic users with complete Grad-CAM and SHAP reports.

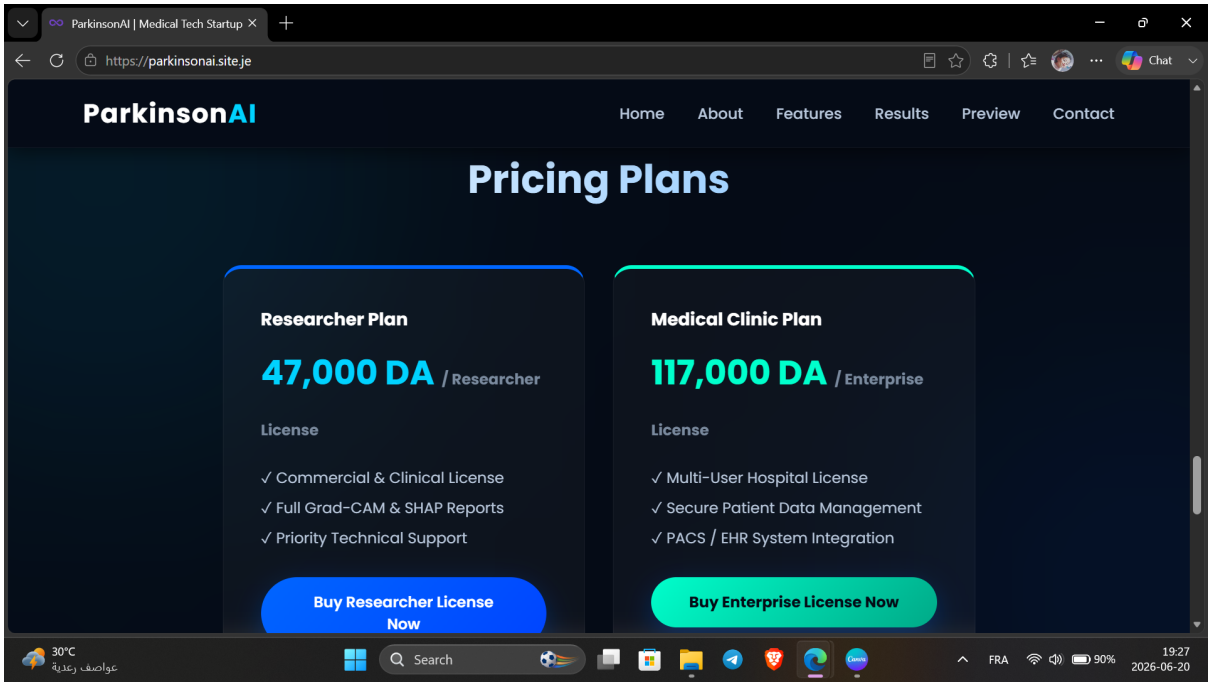


Figure 4.21: ParkinsonAI Web Platform — Pricing Plans

The secure checkout screen is shown in Figure 4.22, where the user enters their card details and billing information before making a payment.

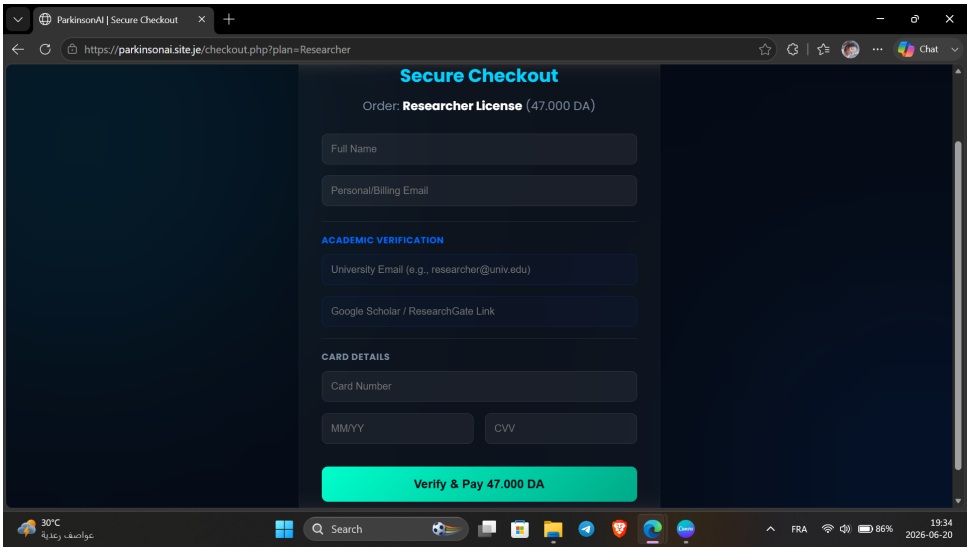


Figure 4.22: ParkinsonAI Web Platform — Secure Checkout page

As seen in Figure 4.23, the platform immediately creates a unique activation key upon successful payment and presents it to the user along with a download button for the Windows installation.

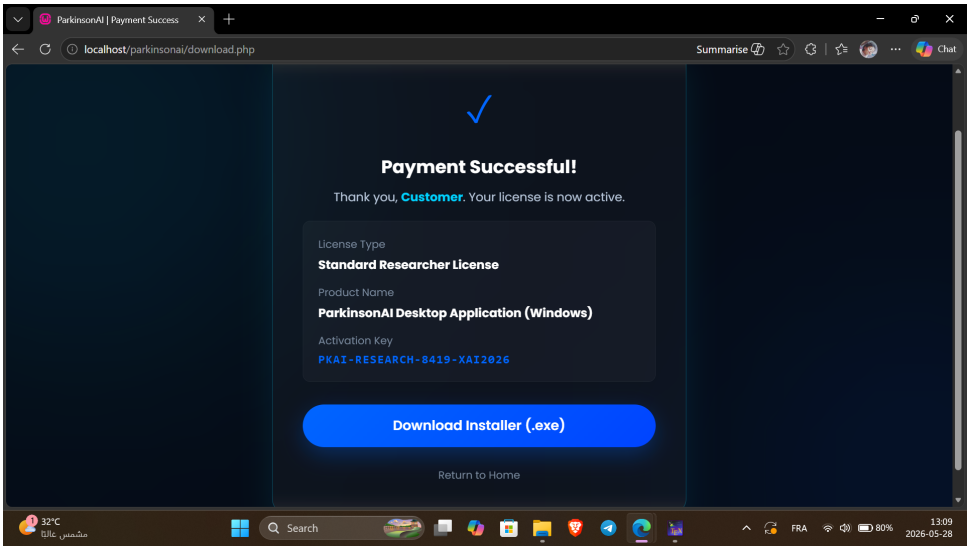


Figure 4.23: ParkinsonAI Web Platform — Payment success and license delivery

The admin dashboard, shown in Figure 4.24, shows real-time statistics such as total revenue, the number of subscriptions for researchers and clinics, and a live sales table with real-time search capabilities that let the administrator filter clients by name, email, or activation key.

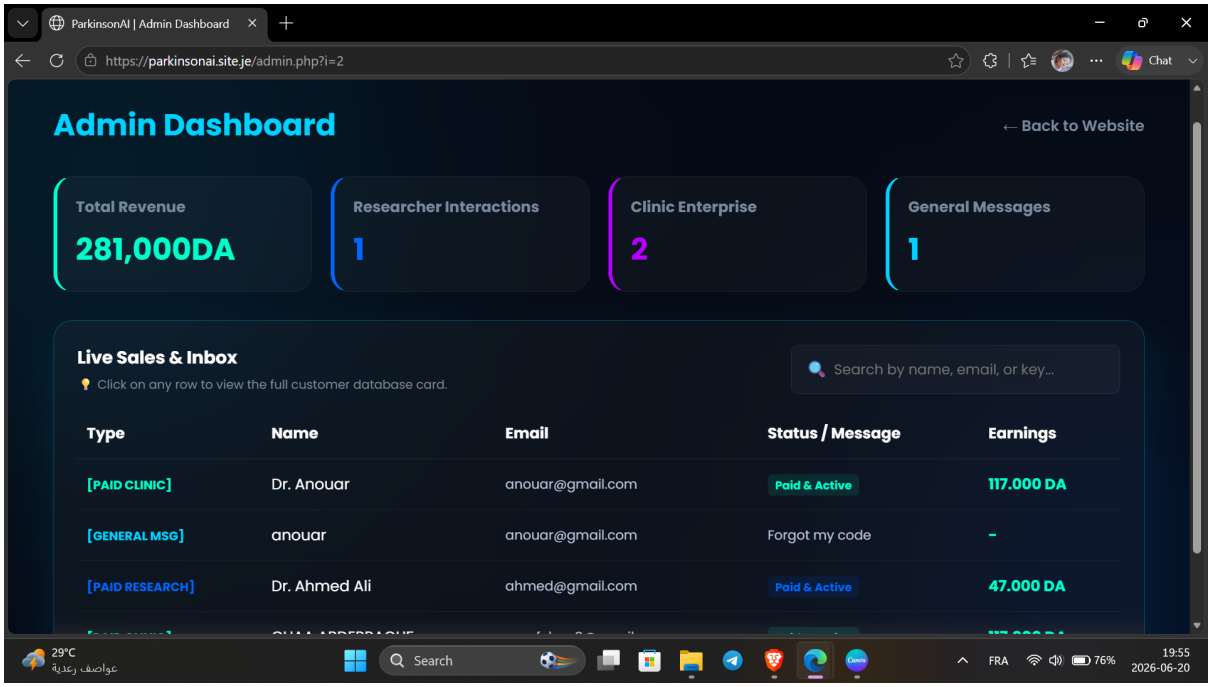


Figure 4.24: ParkinsonAI Web Platform — Admin Dashboard

As seen in Figures 4.25 and 4.26, the platform also has a contact form for order queries and an application preview area with images of the desktop App interface.

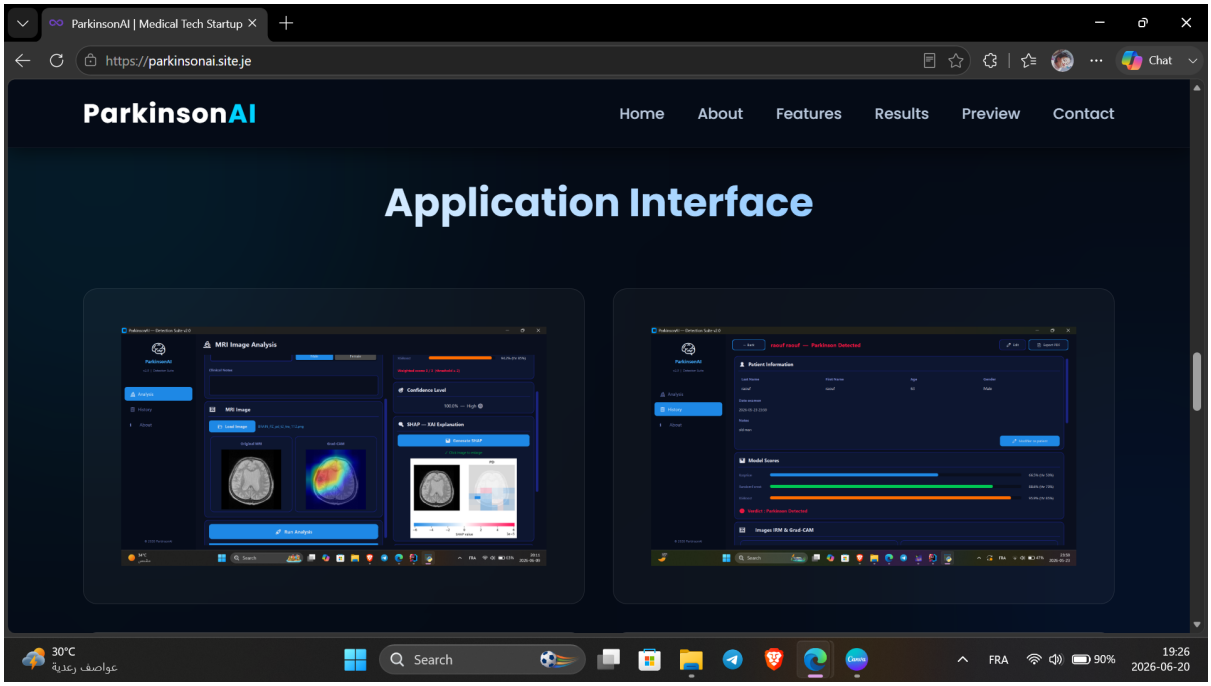


Figure 4.25: ParkinsonAI Web Platform — Application Interface Preview

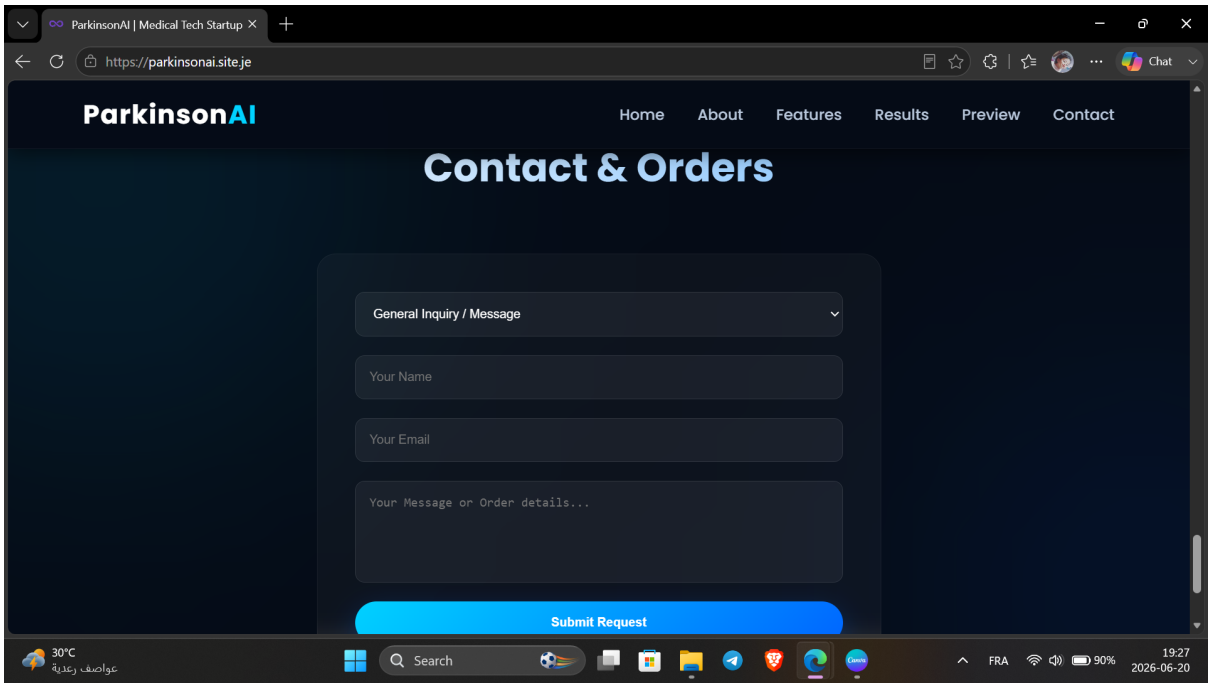


Figure 4.26: ParkinsonAI Web Platform — Contact and Orders page

4.8 Conclusion

The experimental findings of the suggested Parkinson's disease detection method were shown in this chapter. With an accuracy of 84.19%, the Xception CNN outperformed Random Forest (81.40%) and XGBoost (79.13%). The model focuses on therapeutically significant brain areas, especially the Substantia Nigra and Basal Ganglia, according to the XAI analysis employing Grad-CAM, SHAP, and segmentation. Lastly, the ParkinsonAI desktop program and its related commercial online platform were showcased, illustrating the usefulness of the suggested approach for researchers and medical practitioners.

Chapter 5

General Conclusion

5.1 Contributions

This thesis used the capabilities of explainable artificial intelligence, deep learning, and conventional machine learning to propose a comprehensive and integrated system for the automated identification of Parkinson's disease from brain MRI scans. The following is a summary of this work's primary contributions. Initially, a deep learning pipeline based on the Xception convolutional neural network was created and trained using the roughly 44,000 MRI images in the NTUA Parkinson dataset. The model had a significant capacity to differentiate between PD and non-PD instances, with an accuracy of 84.19% on the test set. Second, a hybrid classification strategy was put out in which two traditional machine learning models—Random Forest (81.40%) and XGBoost (79.13%)—are fed by the Xception network, which acts as a feature extractor. A majority vote approach that combines the three models determines the final diagnosis, guaranteeing a more solid and trustworthy forecast. Third, the black-box character of deep learning models was addressed by integrating three explainable AI techniques: SHAP to measure feature relevance, Grad-CAM to show the brain areas impacting the model's judgments, and a segmentation technique to identify aberrant brain regions linked to Parkinson's disease. These methods verified that the model concentrates on clinically significant regions, especially the Basal Ganglia and Substantia Nigra.

Finally, CustomTkinter was used to create ParkinsonAI, a fully working desktop program that integrated the whole detection process with patient management, XAI visualization, medi-

cal report production, and PDF export. Using PHP and JavaScript, a supplementary commercial online platform was also created to guarantee the application's distribution and accessibility to researchers and medical facilities around the globe.

5.2 Limitations

Even though this study produced encouraging results, a number of limitations need to be noted. First, there is a notable class imbalance in the NTUA Parkinson dataset utilized in this study, with PD samples (32,706) significantly outnumbering non-PD samples (10,381). The Random Forest and XGBoost models' poor recall scores for the Non-PD class show that, despite the use of class weighting to address this problem, the models still have a propensity to categorize the majority of samples as PD.

Second, the trained models' capacity to generalize to MRI images obtained from other scanners, procedures, or demographics may be limited because the dataset comes from a single source. One major obstacle to clinical adoption is the lack of multi-center validation. Third, even if the Xception model's accuracy is adequate, it still performs below the standards usually needed for clinical decision support systems. If used in a genuine clinical situation, the rather poor accuracy for the Non-PD class (0.64) suggests a non-negligible proportion of false positives, which might cause needless concern for healthy individuals.

Lastly, the desktop application's immediate practical relevance in clinical processes is limited because it presently runs in an offline environment and does not allow interface with current hospital information systems like PACS or EHR platforms.

5.3 Future Work and Perspectives

In order to overcome the limitations noted and expand the contributions of this thesis, a number of avenues might be investigated in further research.

From a data standpoint, gathering and combining multi-center MRI datasets from various hospitals and scanner manufacturers would greatly enhance the models' capacity for generalization. To solve the class imbalance issue, data augmentation methods created especially for medical imaging and the creation of synthetic data using Generative Adversarial Networks

(GANs) should be researched [11].

In order to possibly enhance classification performance, more sophisticated designs like Vision Transformers (ViT) and hybrid CNN-Transformer models might be investigated. Furthermore, a more thorough diagnostic framework may be offered by multimodal learning techniques that integrate MRI pictures with clinical information, genetic biomarkers, and patient history. From an explainability standpoint, the current Grad-CAM and SHAP analysis might be enhanced by more contemporary XAI approaches like LIME, Integrated Gradients, and attention-based visualization techniques to produce richer and more clinically interpretable explanations [18].

Lastly, from an application standpoint, integrating the system into current PACS and EHR systems and making it available to neurologists globally as a cloud-based online application would be a major step toward actual clinical implementation. To determine the actual clinical efficacy of the suggested method, prospective clinical validation trials including radiologists and neurologists would also be crucial.

References

[1] GBD 2016 Parkinson's Disease Collaborators. Global, regional, and national burden of Parkinson's disease. *The Lancet Neurology*, 17(11):939–953, 2018.

[2] Lorraine V. Kalia and Anthony E. Lang. Parkinson's disease. *The Lancet*, 386(9996): 896–912, 2015.

[3] Bastiaan R. Bloem, Michael S. Okun, and Christine Klein. Parkinson's disease. *The Lancet*, 397(10291):2284–2303, 2021.

[4] Ronald B. Postuma et al. Mds clinical diagnostic criteria for Parkinson's disease. *Movement Disorders*, 30(12):1591–1601, 2015.

[5] A. Jon Stoessl et al. Imaging insights into basal ganglia function. *The Lancet*, 384(9942): 532–544, 2014.

[6] Imtiaz Ahmed et al. Artificial intelligence and machine learning in Parkinson's disease. *Journal of Parkinson's Disease*, 11(1):45–67, 2021.

[7] Mark A. Brown and Richard C. Semelka. *MRI: Basic Principles and Applications*. Wiley-Blackwell, 2014.

[8] Do-Hyung Kwon et al. Seven-tesla mri of the substantia nigra in Parkinson's disease. *Investigative Radiology*, 47(7):427–432, 2012.

[9] Geert Litjens et al. A survey on deep learning in medical image analysis. *Medical Image Analysis*.

[10] Juan Manuel Gorriz et al. Artificial intelligence within the interplay between natural and artificial computation. *Information Fusion*, 59:13–38, 2019.

[11] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.

[12] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521:436–444, 2015.

[13] Yousra Safinez Rahal. Une approche hybride pour la classification des irm de la maladie d'alzheimer : Combinaison de modèles cnn et de classifieurs machine learning avec augmentation de données. Mémoire de master, Université de Bordj Bou Arréridj (BBA), 2025.

[14] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2010.

[15] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1251–1258, 2017.

[16] Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.

[17] Tianqi Chen and Carlos Guestrin. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794, 2016.

[18] Alejandro Barredo Arrieta et al. Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges. *Information Fusion*, 58:82–115, 2020.

[19] Ramprasaath R. Selvaraju et al. Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 618–626, 2017.

[20] Min Lin, Qiang Chen, and Shuicheng Yan. Network in network. *arXiv preprint arXiv:1312.4400*, 2013.

[21] Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, volume 30, pages 4765–4774, 2017.

[22] Lloyd S. Shapley. A value for n-person games. In Harold W. Kuhn and Albert W. Tucker,

editors, *Contributions to the Theory of Games II*, volume 28 of *Annals of Mathematics Studies*, pages 307–317. Princeton University Press, Princeton, NJ, 1953.

[23] Athanasios Tagaris, Dimitrios Kollias, Andreas Stafylopatis, Georgios Tagaris, and Stefanos Kollias. Machine learning for neurodegenerative disorder diagnosis—survey of practices and launch of benchmark dataset. *International Journal on Artificial Intelligence Tools*, 27(03):1850011, 2018.

[24] Nobuyuki Otsu. A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, 9(1):62–66, 1979. doi: 10.1109/TSMC.1979.4310076.

[25] Fabian Pedregosa et al. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[26] Wes McKinney. Data structures for statistical computing in python. In *Proceedings of the 9th Python in Science Conference*, pages 51–56, 2010.

[27] Martin Abadi et al. Tensorflow: Large-scale machine learning on heterogeneous systems, 2015. URL https://www.tensorflow.org.

[28] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.

[29] Tom Fawcett. An introduction to roc analysis. *Pattern Recognition Letters*, 27(8):861–874, 2006.