



OPEN

Deep learning predicts prevalent and incident Parkinson's disease from UK Biobank fundus imaging

Charlie Tran¹, Kai Shen¹, Kang Liu², Akshay Ashok³, Adolfo Ramirez-Zamora⁴, Jinghua Chen⁵, Yulin Li⁶ & Ruogu Fang^{1,7,8}✉

Parkinson's disease is the world's fastest-growing neurological disorder. Research to elucidate the mechanisms of Parkinson's disease and automate diagnostics would greatly improve the treatment of patients with Parkinson's disease. Current diagnostic methods are expensive and have limited availability. Considering the insidious and preclinical onset and progression of the disease, a desirable screening should be diagnostically accurate even before the onset of symptoms to allow medical interventions. We highlight retinal fundus imaging, often termed a window to the brain, as a diagnostic screening modality for Parkinson's disease. We conducted a systematic evaluation of conventional machine learning and deep learning techniques to classify Parkinson's disease from UK Biobank fundus imaging. Our results suggest Parkinson's disease individuals can be differentiated from age and gender-matched healthy subjects with 68% accuracy. This accuracy is maintained when predicting either prevalent or incident Parkinson's disease. Explainability and trustworthiness are enhanced by visual attribution maps of localized biomarkers and quantified metrics of model robustness to data perturbations.

Parkinson's disease (PD) is one of the world's fastest-growing neurological disorders, characterized by progressive impairment in motor control and multiple non-motor symptoms^{1–3}. The manifestation of these symptoms is pathologically characterized by the significant loss of dopaminergic neurons in the substantia nigra⁴. An estimated one million individuals in the United States have PD, leading to nearly \$50 billion a year in economic burden¹. Notably, this financial burden consists not only of direct medical costs but also indirect influences such as necessitated family care and social welfare. The World Health Organization estimates the prevalence of PD has doubled in the last 25 years, while the number of deaths caused by PD increased by over 100% since 2000, largely due to the lack of effective intervention under the rising growth of the elderly population². Innovations to our understanding of the pathology of PD and the development of early diagnostic systems are desired to address the global concerns arising from PD.

Systematic diagnostic evaluation of PD currently struggles due to the lack of early biomarkers and balance between both specificity and sensitivity³. Indeed, cardinal motor symptoms fall within the umbrella of Parkinsonism, while non-motor indicators are symptomatic of numerous neurodegenerative diseases⁵. Risk factors such as age, gender, and environmental toxin exposure are not specific to PD. Differential diagnosis frameworks established by the UK's Parkinson's Disease Society Brain Bank and the International Parkinson and Movement Disorder Society⁶ are the current standards for evaluation. However, these *checklists* require increasing amounts of exclusion and evidence, including response to dopaminergic therapy (levodopa)⁷, the occurrence of dyskinesia⁸, and even DaTscan imaging⁹ to conclude a definitive PD diagnosis. A biological definition necessitating additional testing is being entertained. Moreover, early or atypical PD complicates the observability of cardinal signs, leading to significantly reduced diagnostic accuracy^{10,11}. Thus, current diagnostic indicators lack solid predictive power, straining diagnostic expenses, time, availability, enrollment in therapeutic or disease-modifying trials, and subjectivity.

¹Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL 32611, USA. ²Department of Physics, University of Florida, Gainesville, FL 32661, USA. ³Department of Computer and Information Science and Engineering, University of Florida, Gainesville, FL 32611, USA. ⁴Department of Neurology, University of Florida, Gainesville, FL 32661, USA. ⁵Department of Ophthalmology, University of Florida, Gainesville, FL 32661, USA. ⁶Department of Biostatistics, University of Florida, Gainesville, FL 32661, USA. ⁷J. Crayton Pruitt Family Department of Biomedical Engineering, Herbert Wertheim College of Engineering, University of Florida, 1275 Center Drive, PO Box 116131, Gainesville, FL 32611-6131, USA. ⁸Center for Cognitive Aging and Memory, University of Florida, Gainesville, FL 32611, USA. ✉email: ruogu.fang@ufl.edu

The retina provides a routine, inexpensive, and non-invasive modality for studying brain-related pathological processes of neurodegenerative diseases, often referred to as a *window to the brain*^{12–15}. Dopamine plays a complex role in visual pathway processing, justifying visual dysfunction findings in PD individuals¹⁶. Moreover, substantial evidence has revealed large temporal gaps between the onset of PD and observable symptoms, inspiring the retina as a prodromal PD biomarker⁵. Clinical studies have suggested retinal layer thinning and reductions in microvasculature density in PD patients primarily through optical coherence tomography (OCT) and optical coherence tomography angiography (OCT-A)^{17,18}. Nonetheless, clinical findings concerning both retinal degeneration and disease-specific information (disease duration, disease severity, etc.) are not always consistent^{19,20}, demanding further studies to bolster retinal diagnostic power.

Artificial intelligence (AI) algorithms are efficient diagnostic tools through their ability to identify, localize, and quantify pathological features, as evident from their success in diverse retinal disease tasks^{21,22}. In particular, supervised (labeled) algorithms employed in these tasks can generally be divided into two categories: conventional machine learning algorithms and deep learning models. Conventional machine learning models are known to create (non)linear decision boundaries by rules of inference. On the other hand, deep learning models have risen as powerful models due to their ability to learn meaningful representations and extract subtle features in the training process. Learning retinal biomarkers of PD demands an intricate understanding of structural degeneration of the retinal vasculature, a task unfeasible for even experienced ophthalmologists and neurologists. To address this challenge, we propose the usage of AI algorithms to extract the complex relationships existing within the global and local spatial levels of the retina.

We provide one of the first comprehensive artificial intelligence studies of PD classification from fundus imaging.

Our key objective is realized by systematically profiling the classification performance across different stages of Parkinson's disease progression, namely incident (consistent with pre-symptomatic and/or prodromal) PD and prevalent PD. Improving upon other related works, we maximize the diagnostic capacity of AI algorithms by neglecting the usage of any external quantitative measures or feature selection methods. Finally, we assess diagnostic consistency and robustness through extensive experimentation of both conventional machine learning and deep learning approaches together with a post-hoc spatial feature attribution analysis. Overall, this work enables future research in the development of efficient diagnostic technology and early disease intervention.

Results

Study design and clinical characteristics

This study draws from the UK Biobank, a biomedical database containing over 500,000 participants aged 40–69 from 2006 to 2010²³. As of 2019 October, 175,824 fundus images from 85,848 participants were available along with broad clinical health measures. Most subjects have two retinal photographs (left and right eye) with a minority having an additional follow-up imaging session. From this population, we identified 585 fundus images from 296 unique subjects. Following image quality selection guidelines, we determined 123 fundus images from 84 unique participants. To carry out an unbiased binary classification framework, we matched each PD image according to the subjects' age and gender, that is, a healthy control cohort of 123 fundus images from 84 subjects. This constitutes our binary-labeled *overall* dataset of 246 fundus images (123 PD, 123 HC) from 168 subjects (84 PD, 91 HC). Lastly, we form two subsets of the data corresponding to a *prevalent dataset* of 154 fundus images (77 PD, 77 HC) from 110 subjects (55 PD, 55 HC) and an *incident dataset* of 92 fundus images (46 PD, 46 HC) from 58 subjects (29 PD, 29 HC). Notably, we define the diagnostic gap as the difference between the date of image acquisition minus the date of diagnosis, where a negative value is interpreted as having a PD diagnosis before fundus image acquisition (prevalent PD), and a positive value is interpreted as having a PD diagnosis post fundus image acquisition. The data collection pipeline is summarized in Fig. 1.

Risk factors of Parkinson's disease have been extensively studied^{24–26}, including age, gender, ethnicity, Townsend deprivation indices, alcohol consumption, history of obesity-diabetes, history of stroke, and psychotropic medication usage. Moreover, the effects of Parkinson's Disease have been associated with visual symptoms, from which, we acquire a history of diagnostic eye problems and visual acuity measures. We detail the statistical analyses of subject demographics, visual measures, and covariates of our study population in Table 1.

Model design

Conventional machine learning and deep learning models were systemically studied for their performance in PD diagnosis, a binary classification task. We evaluate the performance of each AI model by five randomized repetitions of stratified five-fold cross validation (at the subject level) based on several classification metrics (AUC, accuracy, PPV, NPV, sensitivity, specificity, and F1-score). Given the insufficient amount of longitudinal data, we treat each image separately as a sample without data leakage, rather than performing a longitudinal analysis. The conventional machine learning models provide a classification performance baseline, from which we utilize Logistic Regression, Elastic-Net, Linear SVM, and Radial Basis Function SVM kernels. On the other hand, we evaluate the performance of popular deep learning frameworks including AlexNet, VGG-16, GoogleNet, Inception-V3, and ResNet-50. We follow traditional guidelines such as image normalization to the training data for our machine learning models, ImageNet normalization, spatial augmentations, and early stopping for our deep learning models. The detailed guidelines are outlined in “[Methods](#)” (“*Artificial Intelligence and Model Training*”).

Model performance in overall PD groups

The best deep learning model was AlexNet with an average AUC of 0.77 (95% CI 0.74–0.81), accuracy of 0.68 (95% CI 0.65–0.72), PPV of 0.69 (95% CI 0.64–0.74), NPV of 0.79 (95% CI 0.72–0.86), sensitivity of 0.76 (95% CI 0.66–0.87), specificity of 0.60 (95% CI 0.51–0.70), and F1-score of 0.68 (95% CI 0.62–0.75). The best conventional

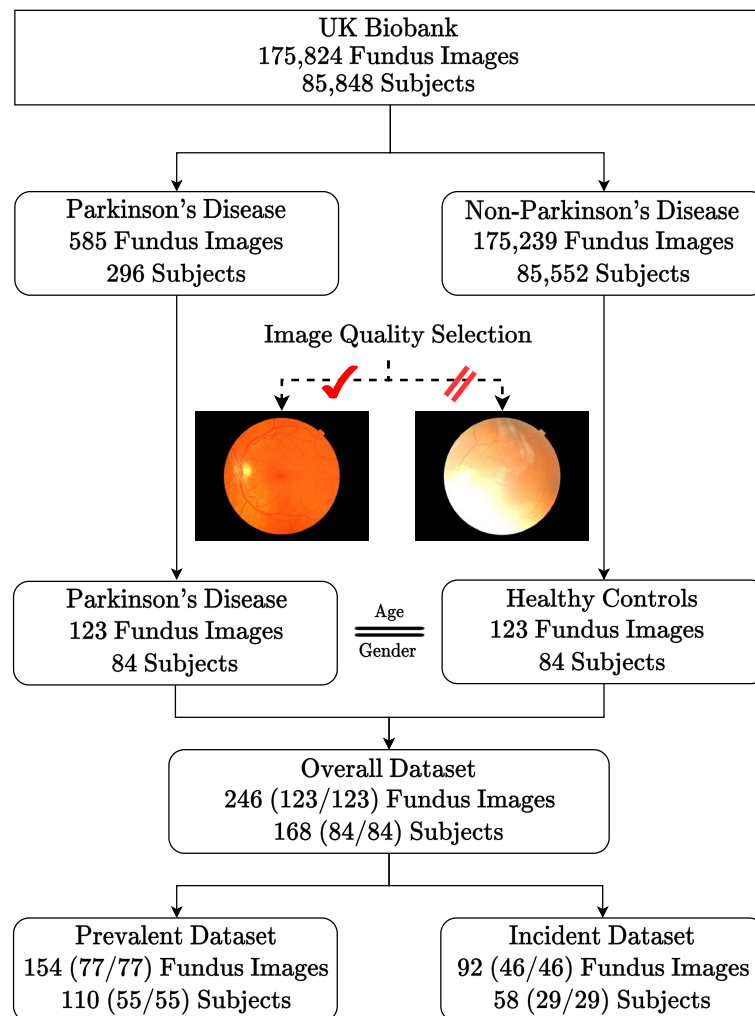


Figure 1. Data collection pipeline from the UK Biobank. Instances in parentheses represent an equal balance of Parkinson's disease and healthy control subjects. Multiple quality selection phases were used as additional inclusion criteria into our dataset arising from AutoMorph and manual image grading. In total, we have the *overall* dataset of PD subjects matched with age and gender-matched healthy controls, and two subsets corresponding to *prevalent* and *incident* subjects.

machine learning model was the RBF support vector machine with an optimal AUC of 0.71 (95% CI 0.69–0.74), accuracy of 0.67 (95% CI 0.65–0.70), PPV of 0.65 (0.63–0.67), NPV of 0.71 (0.68–0.74), sensitivity of 0.76 (95% CI 0.72–0.80), specificity of 0.58 (95% CI 0.54–0.63), and F1-score of 0.70 (95% CI 0.68–0.72). The VGG-16, GoogleNet, SVM (Linear), Logistic Regression, and ElasticNet perform similarly, while Inception-V3 and ResNet-50 perform significantly worse (Fig. 2 and Table 2).

Model performance in prevalent PD groups

The best performing deep learning was the AlexNet model with an average AUC of 0.73 (95% CI 0.68–0.77), accuracy of 0.65 (95% CI 0.62–0.68), PPV of 0.62 (95% CI 0.55–0.68), NPV of 0.69 (95% CI 0.61–0.77), sensitivity of 0.74 (95% CI 0.65–0.83), specificity of 0.56 (95% CI 0.47–0.65), and F1-score of 0.66 (95% CI 0.60–0.72). The best performing conventional machine learning model was the SVM (Linear) with an AUC of 0.73 (95% CI 0.69–0.77), accuracy of 0.67 (95% CI 0.64–0.70), PPV of 0.67 (95% CI 0.63–0.72), NPV of 0.68 (95% CI 0.64–0.71), sensitivity of 0.69 (95% CI 0.64–0.73), specificity of 0.65 (95% CI 0.58–0.71), and F1-score of 0.67 (95% CI 0.64–0.70). Note that in contrast to the “Overall PD Groups”, the SVM (Linear) outperformed the SVM (RBF). In general, the performance of the artificial intelligence models was relatively slight upon the transition between “Overall” to “Prevalent” PD groups. Detailed comparisons can be found in Fig. 2 and Table 2.

Model performance in incident PD groups

The best-performing deep learning model was the AlexNet model with an average AUC of 0.68 (95% CI 0.60–0.75), accuracy of 0.60% (95% CI 0.54–0.66), PPV of 0.49 (95% CI 0.39–0.60), NPV of 0.66 (95% CI 0.56–0.77), sensitivity of 0.64 (95% CI 0.48–0.79), specificity of 0.56 (95% CI 0.45–0.68) and F1-score of 0.54

Baseline characteristics	PD group	HC group	p-value	Prevalent	Incident	p-value
N	84	84	–	55	29	–
Diagnostic gap, mean (SD), years				– 5.7 (4.5)	4.8 (1.4)	–
Age, mean (SD), years	61.8 (5.9)	61.8 (5.9)	0.97	61.4 (5.8)	62.7 (6.1)	0.33
Gender, no. (%)			1.0			0.97
Male	49 (58.3)	49 (58.3)		32 (58.2)	17 (58.6)	
Female	35 (41.7)	35 (41.7)		23 (41.8)	12 (41.4)	
Ethnicity, no. (%)			0.56			0.64
White	82 (97.6)	83 (98.2)		54 (98.2)	28 (96.6)	
Others	2 (2.4)	1 (1.2)		1 (1.8)	1 (3.4)	
Eye problems, no. (%)			0.15			0.94
Yes	9 (10.7)	4.0 (4.8)		6 (10.9)	3 (10.3)	
No	75 (89.3)	80 (95.2)		49 (89.1)	26 (89.7)	
Diabetes-related	0	1		0	0	
Glaucoma	2	0		0	2	
Injury/trauma	1	0		1	0	
Cataracts	5	2		4	1	
Macular degeneration	0	1		0	1	
Other	1	0		1	0	
Visual acuity, mean (SD)	0.28 (0.38)	0.04 (0.37)	*	0.28 (0.37)	0.29 (0.41)	0.92
Townsend index, mean (SD)	– 1.48 (2.95)	– 2.68 (2.17)	0.73	– 1.59 (3.09)	– 1.27 (2.68)	0.64
Smoking status, no. (%)			0.62			0.08
Former/current	30 (35.7)	27 (32.1)		16 (29.1)	14 (48.3)	
Non-smoker	54 (64.3)	57 (67.9)		39 (70.9)	15 (51.7)	
Drinking status, no. (%)			0.73			0.64
Former/current	80 (95.2)	79 (94.0)		52 (94.5)	30 (97.8)	
Non-smoker	4 (4.8)	5 (6.0)		3 (5.5)	1 (3.2)	
Obesity-diabetes, no. (%)			0.46			0.14
Yes	21 (25.0)	17 (20.2)		11 (20.0)	10 (34.5)	
No	63 (75.0)	67 (79.8)		44 (80.0)	19 (65.5)	
History of stroke, no. (%)			0.47			0.23
Yes	3 (3.6)	5 (6.0)		1 (1.8)	2 (6.9)	
No	81 (96.4)	79 (94.0)		54 (98.2)	27 (93.1)	
Psychotropic medication, no. (%)			0.32			–
Yes	0 (0)	1 (1.2)		0 (0)	0 (0)	
No	84 (100)	83 (98.8)		55 (100)	29 (100)	

Table 1. Baseline characteristics of the study populations. P-values conducted on continuous data are computed by the Student's t-test while categorical (binary) variables are computed by Pearson's Chi-squared test. *Indicates statistically significant ($p < 0.01$).

(95% CI 0.42–0.67). The performance of the AlexNet and VGG-16 remains similar while other deep learning models decline with less than 65% accuracy. The best-performing conventional machine learning model was the SVM (RBF) with an AUC of 0.64 (95% CI 0.59–0.70), accuracy of 0.62 (95% CI 0.58–0.67), PPV of 0.63 (95% CI 0.58–0.68), NPV of 0.63 (95% CI 0.58–0.69), sensitivity of 0.64 (95% CI 0.59–0.69), specificity of 0.60 (95% CI 0.53–0.67), and F1-score of 0.63 (95% CI 0.59–0.67). In general, the performance of the machine learning classifiers decreased significantly relative to the overall and prevalent groups. A summarized completion of the results can be found in Fig. 2 and Table 2.

Visualization and explainability

Qualitative explainability is visualized through guided backpropagation on our deep learning models, revealing the models were able to correctly distinguish subtle features on the retinal anatomy (Fig. 3). To highlight that the extracted features are consistent with those recognized as retinal biomarkers of neurodegeneration in PD, we use the AutoMorph²⁷ deep learning segmentation module to generate a map of important retinal structures, namely the arteries, veins, optic cup, and optic disc, along with a manually marked fovea. The optic cup, optic disc, and fovea have been shown to flag progression in PD via observable structural changes in size, thickness, and other features^{28–30}. As such, the overlay of the guided backpropagation map and segmentation map yields a visual comparison of the consistency between features used for classification and potential retinal biomarkers of PD. Moreover, we reinforce the qualitative explainability by quantifying the robustness of our deep learning

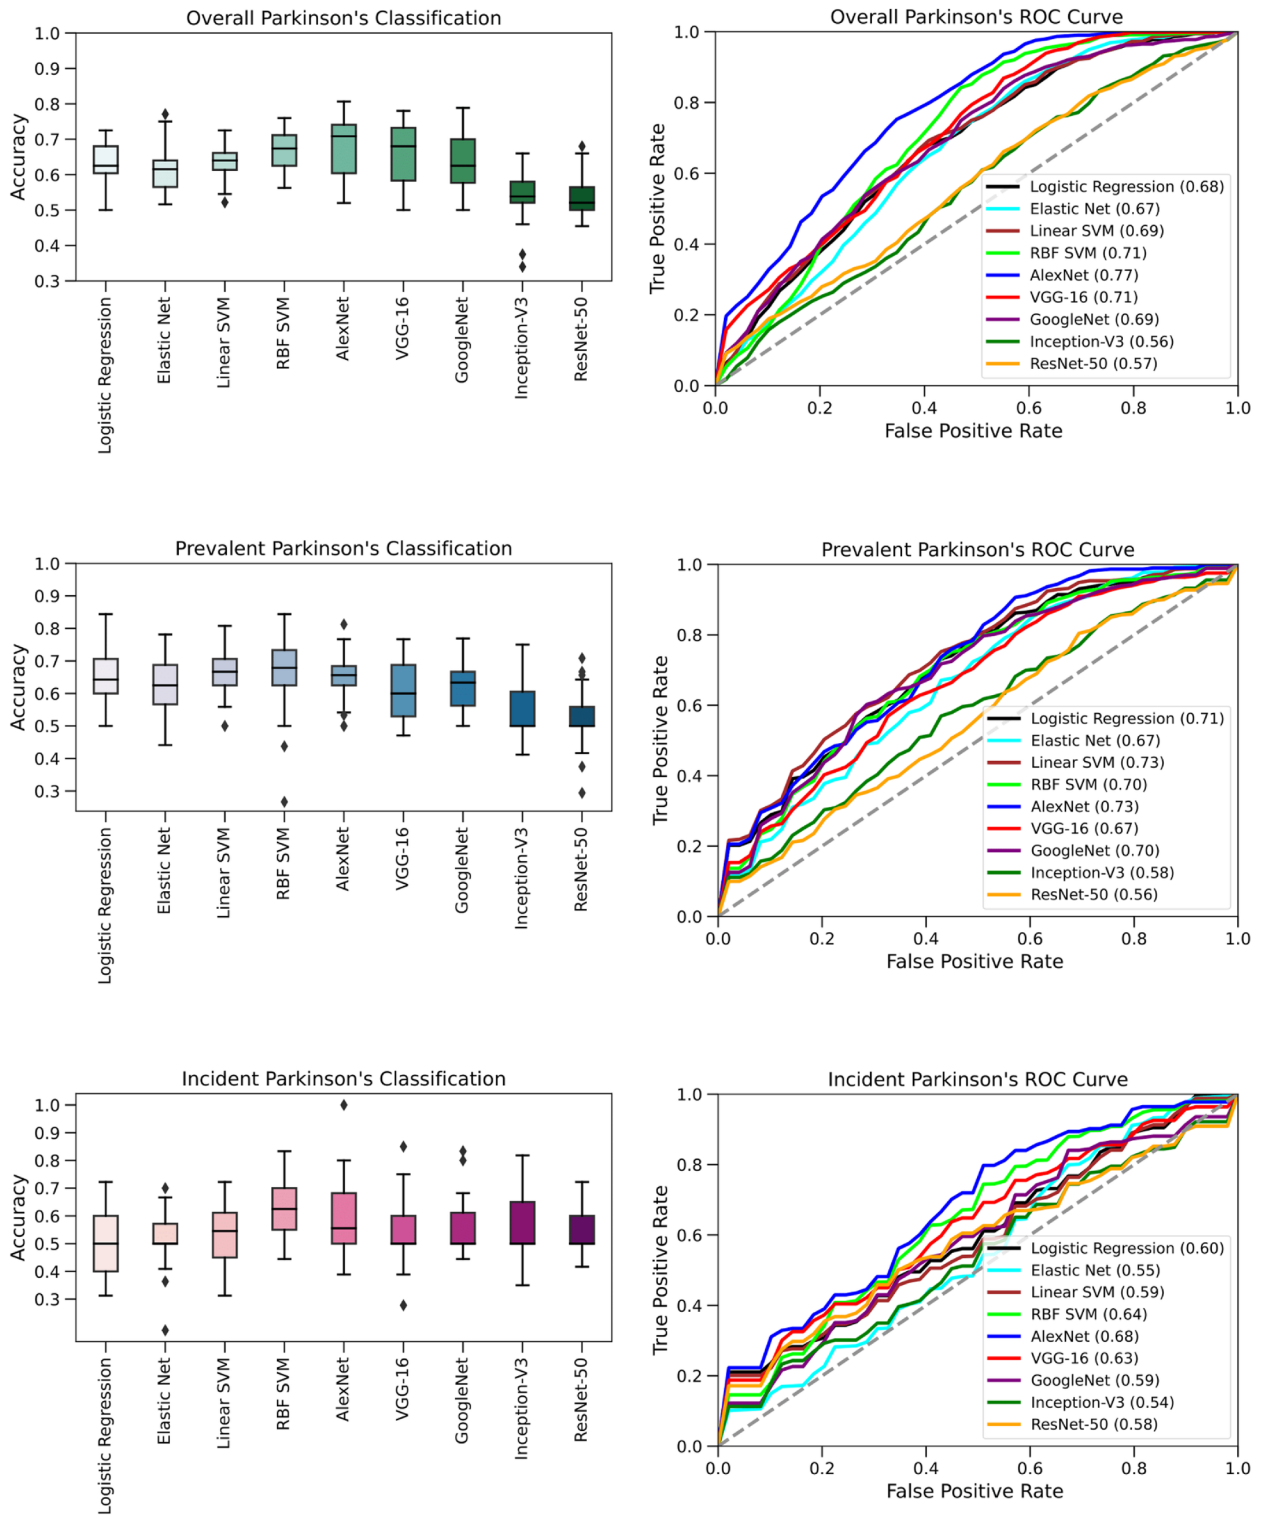


Figure 2. Boxplots and ROC curves of the Parkinson's Disease Classification Models. The models are evaluated over five randomized repetitions of the five-fold stratified cross-validation protocol. The AUC scores are enlisted in the legend.

models to data perturbations through an infidelity and sensitivity analysis. The definitions of the infidelity and sensitivity measures are outlined in “Methods” (“Explainability Evaluation”). In this step, we run the guided backpropagation attribution step across $N = 50$ perturbations at test time with a noise distribution drawn from a normal distribution $N \sim (0, 0.01^2)$. The results of the infidelity and sensitivity analysis are averaged across the repeated cross validation protocol and summarized in Fig. 4. Our empirical results suggest that AlexNet is the most robust to data perturbations, while simultaneously the most accurate according to lower infidelity and

Model	AUC	ACC	PPV	NPV	SENS	SPEC	F1
Overall (n = 246 fundus images 123 PD, 123 HC)							
Logistic regression	0.68 (0.65, 0.71)	0.63 (0.61, 0.66)	0.63 (0.6, 0.66)	0.64 (0.61, 0.67)	0.66 (0.62, 0.70)	0.61 (0.55, 0.66)	0.64 (0.62, 0.66)
Elastic Net	0.67 (0.64, 0.69)	0.61 (0.58, 0.64)	0.60 (0.57, 0.62)	0.63 (0.60, 0.67)	0.68 (0.64, 0.72)	0.54 (0.49, 0.58)	0.64 (0.61, 0.66)
SVM (linear)	0.69 (0.66, 0.71)	0.63 (0.61, 0.66)	0.64 (0.61, 0.66)	0.64 (0.62, 0.66)	0.65 (0.61, 0.68)	0.62 (0.57, 0.67)	0.64 (0.61, 0.66)
SVM (RBF)	0.71 (0.69, 0.74)	0.67 (0.65, 0.70)	0.65 (0.63, 0.67)	0.71 (0.68, 0.74)	0.76 (0.72, 0.80)	0.58 (0.54, 0.63)	0.70 (0.68, 0.72)
Alexnet	0.77 (0.74, 0.81)	0.68 (0.65, 0.72)	0.69 (0.64, 0.74)	0.79 (0.72, 0.86)	0.76 (0.66, 0.87)	0.60 (0.51, 0.7)	0.68 (0.62, 0.75)
VGG-16	0.71 (0.67, 0.75)	0.66 (0.62, 0.69)	0.61 (0.58, 0.64)	0.81 (0.71, 0.91)	0.85 (0.76, 0.95)	0.46 (0.37, 0.54)	0.70 (0.64, 0.76)
GoogleNet	0.69 (0.65, 0.73)	0.63 (0.60, 0.66)	0.62 (0.58, 0.65)	0.70 (0.62, 0.79)	0.77 (0.69, 0.85)	0.49 (0.41, 0.58)	0.67 (0.63, 0.71)
Inception-V3	0.56 (0.52, 0.60)	0.54 (0.51, 0.57)	0.54 (0.49, 0.60)	0.59 (0.52, 0.66)	0.68 (0.57, 0.80)	0.40 (0.29, 0.50)	0.57 (0.50, 0.64)
ResNet-50	0.57 (0.53, 0.62)	0.54 (0.51, 0.57)	0.52 (0.40, 0.63)	0.55 (0.45, 0.64)	0.52 (0.37, 0.68)	0.55 (0.41, 0.70)	0.45 (0.34, 0.56)
Prevalent (n = 146 fundus images 73 PD, 73 HC)							
Logistic regression	0.71 (0.67, 0.76)	0.65 (0.62, 0.69)	0.66 (0.61, 0.70)	0.66 (0.63, 0.70)	0.68 (0.63, 0.73)	0.63 (0.56, 0.69)	0.66 (0.62, 0.70)
Elastic Net	0.67 (0.62, 0.71)	0.62 (0.59, 0.66)	0.62 (0.57, 0.66)	0.63 (0.56, 0.70)	0.71 (0.65, 0.78)	0.53 (0.45, 0.61)	0.65 (0.61, 0.69)
SVM (Linear)	0.73 (0.69, 0.77)	0.67 (0.64, 0.70)	0.67 (0.63, 0.72)	0.68 (0.64, 0.71)	0.69 (0.64, 0.73)	0.65 (0.58, 0.71)	0.67 (0.64, 0.70)
SVM (RBF)	0.70 (0.64, 0.76)	0.66 (0.60, 0.71)	0.64 (0.57, 0.71)	0.69 (0.63, 0.75)	0.70 (0.61, 0.78)	0.62 (0.55, 0.68)	0.66 (0.59, 0.73)
Alexnet	0.73 (0.68, 0.77)	0.65 (0.62, 0.68)	0.62 (0.55, 0.68)	0.69 (0.61, 0.77)	0.74 (0.65, 0.83)	0.56 (0.47, 0.65)	0.66 (0.60, 0.72)
VGG-16	0.67 (0.61, 0.73)	0.61 (0.57, 0.65)	0.58 (0.51, 0.64)	0.63 (0.55, 0.71)	0.66 (0.55, 0.76)	0.56 (0.47, 0.65)	0.60 (0.52, 0.67)
GoogleNet	0.70 (0.66, 0.74)	0.62 (0.59, 0.66)	0.60 (0.54, 0.66)	0.68 (0.63, 0.74)	0.69 (0.60, 0.79)	0.55 (0.46, 0.65)	0.63 (0.56, 0.69)
Inception-V3	0.58 (0.52, 0.64)	0.55 (0.51, 0.59)	0.49 (0.40, 0.58)	0.52 (0.44, 0.60)	0.55 (0.42, 0.68)	0.55 (0.43, 0.67)	0.50 (0.40, 0.59)
ResNet-50	0.56 (0.50, 0.62)	0.52 (0.48, 0.56)	0.49 (0.42, 0.57)	0.49 (0.39, 0.59)	0.51 (0.38, 0.64)	0.53 (0.39, 0.67)	0.46 (0.38, 0.55)
Incident (n = 100 fundus images 50 PD, 50 HC)							
Logistic regression	0.60 (0.54, 0.66)	0.52 (0.47, 0.56)	0.53 (0.48, 0.57)	0.50 (0.44, 0.56)	0.56 (0.52, 0.61)	0.47 (0.39, 0.55)	0.54 (0.50, 0.58)
Elastic Net	0.55 (0.50, 0.61)	0.52 (0.48, 0.56)	0.52 (0.48, 0.56)	0.48 (0.40, 0.56)	0.63 (0.55, 0.71)	0.41 (0.32, 0.49)	0.56 (0.51, 0.61)
SVM (linear)	0.59 (0.53, 0.65)	0.53 (0.48, 0.58)	0.55 (0.50, 0.59)	0.52 (0.46, 0.58)	0.58 (0.52, 0.63)	0.48 (0.40, 0.57)	0.55 (0.51, 0.59)
SVM (RBF)	0.64 (0.59, 0.70)	0.62 (0.58, 0.67)	0.63 (0.58, 0.68)	0.63 (0.58, 0.67)	0.64 (0.59, 0.69)	0.60 (0.53, 0.67)	0.63 (0.59, 0.67)
Alexnet	0.68 (0.60, 0.75)	0.60 (0.54, 0.66)	0.49 (0.39, 0.60)	0.66 (0.56, 0.77)	0.64 (0.48, 0.79)	0.56 (0.45, 0.68)	0.54 (0.42, 0.67)
VGG-16	0.63 (0.54, 0.71)	0.55 (0.50, 0.60)	0.54 (0.47, 0.61)	0.45 (0.31, 0.60)	0.69 (0.57, 0.82)	0.40 (0.26, 0.55)	0.57 (0.50, 0.65)
GoogleNet	0.59 (0.52, 0.66)	0.56 (0.52, 0.61)	0.49 (0.41, 0.58)	0.57 (0.47, 0.67)	0.63 (0.51, 0.76)	0.49 (0.38, 0.60)	0.55 (0.45, 0.64)
Inception-V3	0.54 (0.47, 0.61)	0.56 (0.51, 0.61)	0.56 (0.51, 0.62)	0.50 (0.39, 0.61)	0.64 (0.54, 0.75)	0.48 (0.36, 0.60)	0.57 (0.51, 0.64)
ResNet-50	0.58 (0.50, 0.66)	0.55 (0.52, 0.58)	0.53 (0.43, 0.63)	0.42 (0.30, 0.55)	0.63 (0.48, 0.77)	0.47 (0.31, 0.64)	0.52 (0.43, 0.62)

Table 2. Classification results of Parkinson disease datasets. For each model, the mean average for each performance metric with their 95% confidence interval is provided over 5 randomized repetitions of fivefold stratified cross-validation. Highlighted in bold represents the best model and highlighted in italics is the best performing conventional machine learning model.

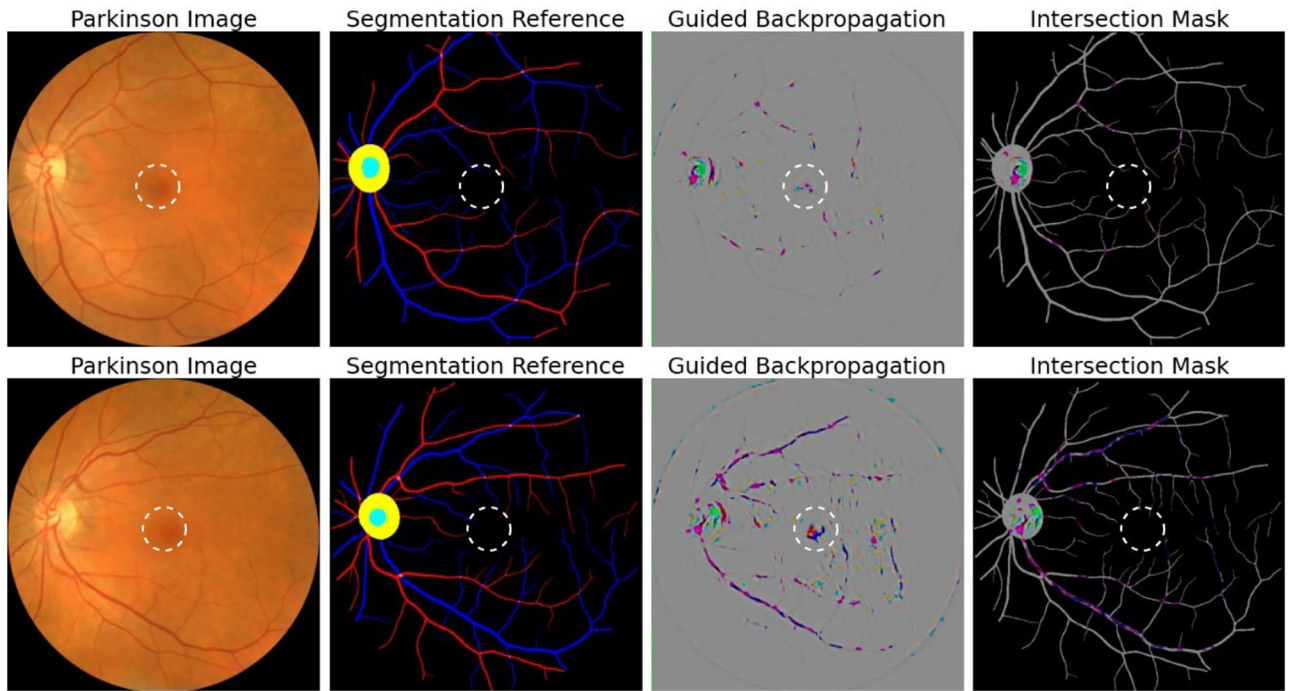


Figure 3. Attribution correspondence of retinal features. In the first column, an artery-vein (red and blue, respectively) map is combined with the optic cup (teal) and optic disc (yellow) generated from the AutoMorph deep learning segmentation module. A white dashed line is shown as an estimate for the foveal region. In the third column, a predicted attribution map is generated using the guided backpropagation algorithm on top of the AlexNet model. The intersection of the salient features with the segmentation is shown in the last column. The images represent the left (top) and right (bottom) eyes from the same subject, demonstrating distinct feature distributions for prediction.

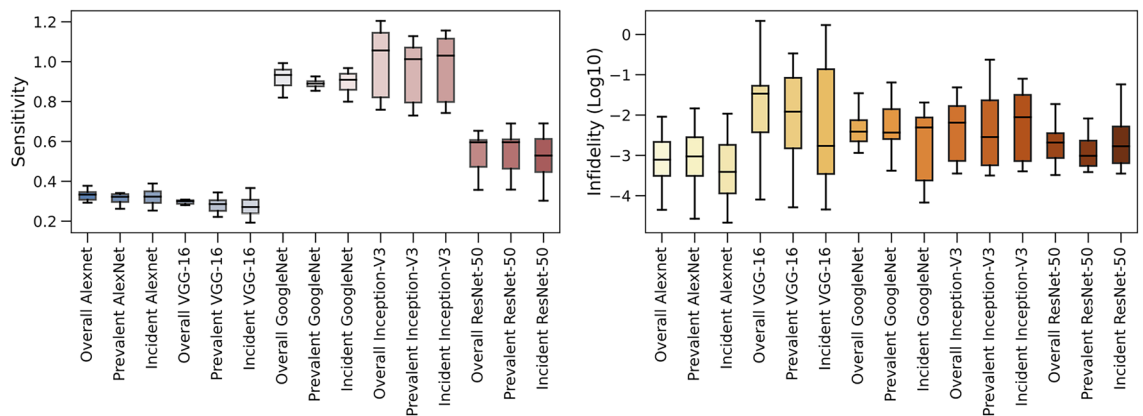


Figure 4. Explanation of sensitivity and infidelity comparison among different models. The logarithm of the infidelity score was applied for visualization due to the large range of scores.

sensitivity measures. Notably, the VGG-16 has similar robustness measures as AlexNet with less classification performance, possibly owing to the large model parameter complexity. Predictions made by other deep learning architectures including GoogLeNet, Inception-V3, and ResNet-50 were largely influenced by data perturbations.

Feature engineering analysis

Feature engineering approaches are generally devised to enhance AI outcomes, in particular, conventional machine learning models, due to the large feature space complexity. On the other hand, deep learning models by design can achieve performance with minimal amounts of data pre-preprocessing. We explore the performance of feature engineering in conventional machine learning models with two changes to the input fundus images: (1) gray-scale color conversion (reducing the dependence of color), and (2) vessel-segmentation via AutoMorph (emphasizing the retinal vasculature). The approach adopted in the latter is also a useful comparison with that of Tian et al. which utilized the technique for Alzheimer’s Disease detection. The results of this approach are

demonstrated in Supplementary Table 4. Conventional machine learning models are demonstrated to improve by simple color conversion, while declining using the retinal vasculature. The reduction in performance using the vessel segmentation algorithm hints that essential diagnostic features exist across different regions of the eye (e.g. the optic cup and fovea).

Model covariate analysis—diagnostic gap and gender

We investigate the influence of Parkinson's disease progression on the model performance. The Parkinson's disease progression can be expressed by the diagnostic gap and thereby a proxy measure of disease severity. Treating each image independently, we divided our Parkinson's subjects into four quartiles according to the diagnostic time (years) and examined the model performance compiled over our repeated five-fold cross validation in our best model, AlexNet. Sensitivity measures in the Parkinson's group did not exhibit monotonic relationships based on the diagnostic gap or consistencies in dataset subtypes (see Supplementary Table 1). One such explanation for this result is discrepancies in the diagnostic gap, that is, the dates of diagnosis do not co-align with the true dates of disease acquisition as Parkinson's disease is a progressive disease. Otherwise, a claim could be made that the model performance is not contingent strictly on the diagnostic gap, but rather, an individual basis of existing imaging biomarkers.

Gender is a known risk factor of Parkinson's disease^{31,32}. We investigated a potential bias in our deep learning models due to gender differences in the retina. We categorized the model performances of our AlexNet model compiled over all subjects, Parkinson's-specific, and healthy-control specific data, and conducted a series of Chi-Square Test of Independence (see Supplementary Table 2). In cases where the frequency of observations was found to be less than or equal to five, we conducted a Fisher's Exact test. We discover no statistical significance ($p < 0.05$) for any experiment or data-subtype.

Discussion

This work demonstrates that deep neural networks can be trained to detect Parkinson's disease in retinal fundus images with decent performance. Our model can predict the incidence of Parkinson's disease ahead of formal diagnosis at demonstrated sensitivity levels of 80.0% from 0 to 3.93 years, 80.0% from 3.93 to 5.07 years, 93.33% from 5.07 to 5.57 years, and 81.67% from 5.57 to 7.38 years. These results indicate a potential pathway for early disease intervention. Automated deep neural networks show strong promises to assist and complement ophthalmologists in terms of biomarker identification and high-throughput evaluation.

Artificial intelligence evaluation of Parkinson's disease through the retina has been rarely applied. Hu et al.³³ trained a deep learning model to evaluate the retinal age gap as one predictive marker for incident Parkinson's Disease using fundus images from the UK Biobank, showcasing statistical significance and predictive AUC of 0.71. Nunes et al.³⁴ used optical coherence tomography data to compute retinal texture markers and trained a deep learning model with a median sensitivity of 88.7%, 79.5%, and 77.8% concerning healthy controls, Parkinson's disease, and Alzheimer's disease. However, all these works did not provide a comprehensive comparison of conventional machine learning and deep learning methods on this problem and lacked insights into the explainability of their models. We extend upon these works by treating the entire fundus image as a diagnostic modality, and comprehensively evaluate a broad spectrum of conventional machine learning and deep learning methods, as well as shedding light on the explainability in both image space and on the algorithm level. Our work will lay a solid foundation for future exploration in this direction and serve as a reference for algorithm selection in terms of both performance and explainability. Related works have been accomplished focusing specifically on Alzheimer's disease, e.g., Tian et al.³⁵ and Wisely et al.³⁶ but not Parkinson's disease. Clinical studies in the field have yielded statistical differences in the retinal layers between PD and HC subjects, lacking evidence for diagnostic power. Further work is necessitated in the field of deep learning to build stronger classification performance and understanding of retinal biomarkers. In the future, a multi-modal model utilizing optical coherence tomography, fundus autofluorescence, and/or electronic health records is a considerable direction for Parkinson's disease analysis.

This study has some limitations. First, the size of our dataset could be enlarged to further capture the wide presentations of Parkinson's disease. Moreover, the data is derived from the UK population and therefore future studies are needed to evaluate whether these models can generalize to other populations. So far, public datasets containing both Parkinson's disease subjects and fundus images (as well as patient health records) apart from the UK Biobank are not available, thus new datasets including both retinal images and PD diagnosis in a larger population will be helpful. A deeper study would hope to investigate the severity of Parkinson's disease, e.g. by the MDS-Unified Parkinson's Disease Rating Scale (MDS-UPDRS), wherein the severity was (weakly) substituted by the diagnostic gap in this work. Furthermore, this current research has been restricted to Parkinson's disease, and it remains questionable whether different eye diseases or neurodegenerative diseases (e.g., Alzheimer's Disease) share identical biomarkers or degeneration patterns. The next major research question is whether such model explanations are consistent and/or able to guide the grading of ophthalmologists, which is a major goal of clinical translational research. This matter is further complicated as the visual biomarkers for Parkinson's disease are less well-understood than common eye diseases such as glaucoma. Prospective assessments of retina imaging coupled with biological details and clinical phenotyping are needed to provide insight into the use and implementation of these techniques. These limitations necessitate future work to ensure the trustworthiness of artificial intelligence models in a clinical setting.

Deep learning models outperformed conventional machine learning models to accurately predict Parkinson's disease from retinal fundus images. We demonstrate deep learning models can nearly equally diagnose both prevalent and incident PD subjects with robustness to image perturbations, paving the way for early treatment

and intervention. Further studies are warranted to verify the consistency of Parkinson's disease evaluation, to enhance our understanding of retinal biomarkers, and to incorporate automated models into clinical settings.

Methods

UK Biobank participants

The UK Biobank (UKB) is one of the largest biomedical databases, recruiting over 500,000 individuals aged 40–69 years old at baseline throughout assessment centers in the United Kingdom in 2006–2010. The methods by which this data was acquired have been described elsewhere²³. Diverse patient health records include demographic, genetic, lifestyle, and health information. Comprehensive physical examinations as well as ophthalmic examinations were conducted for further analysis. Health-related events were determined using data linkage to the Health Episode Statistics (HES), Scottish Morbidity Record (SMR01), Patient Episode Database for Wales (PEDW), and death registers.

Parkinson's disease and definitions

Parkinson's disease was determined by hospital administration data in the United Kingdom, national death register data, and self-reported data. We consider *prevalent* PD subjects diagnosed prior to baseline assessment and *incident* PD subjects diagnosed following baseline assessment. Prevalent PD subjects were labeled according to hospital admission electronic health records based on the International Classification of Diseases (ICD9, ICD10) codes or self-reports. Incident PD subjects were labeled according to either the ICD codes or the death registry. The earliest recorded diagnostic dates take priority in case of multiple records. If PD was recorded in the death register only (diagnosed post-mortem), the date of death is used for the date of diagnosis. We define the diagnostic gap as the difference between the date of image acquisition minus the date of diagnosis, where a negative value is interpreted as having a PD diagnosis prior to fundus image acquisition (prevalent PD), and a positive value is interpreted as having a PD diagnosis post fundus image acquisition (incident PD). We acquire our PD labels according to the UKB Field 42032. Further details may be inquired upon from the UKB documentation³⁷.

Ophthalmic measures

In the UKB eye and vision consortium, the ophthalmic assessment included (1) questionnaires of past ophthalmic and family history, (2) quantitative measures of visual acuity, refractive error, and keratometry, and (3) imaging acquisition including spectral domain optical coherence tomography (SD-OCT) of the macula and a disc-macula fundus photograph. In our study, we acquire diagnostic fields of eye problems (glaucoma, cataracts, diabetes-related, injury/trauma, macular, degeneration, etc.), visual acuity measured as the logarithm of the minimum angle of resolution (LogMar), and fundus photographs. Fundus photographs were acquired using a Topcon 3D OCT-1000 Mark II system. The system has a 45° field angle, scanning range of 6 mm × 6 mm centered on the fovea, acquisition speed of 18,000 A-scans per second, and 6 μm axial resolution. The details of the eye and vision consortium have been described in other studies³⁸.

Study population and summary statistics

A total of 175,824 fundus images from 85,848 subjects were discovered in the UKB (as of 2019 October, UKB). Among this population, we found 585 fundus images from 296 subjects with PD. Image quality selection was held in multiple phases: (1) the deep learning image quality selection module of AutoMorph²⁷ pretrained on Eye-PACS-Q³⁹, (2) subjectivity to the quality of the vessel and optic cup and disc AutoMorph segmentation, and (3) manual grading borderline images according to external guidelines²¹, including artifacts, clarity, and field definition defects. Each phase is held in sequence with the additional rounds held to justify borderline candidates for inclusion or exclusion into the dataset. In these guidelines, an image quality score is determined from the total of artifacts (– 10 to 0), clarity (0 to + 10), and field definition (0 to + 10) where an optimal score is + 20 and a score less than or equal to 12 is rejectable. The artifacts component evaluates the broad proportion of visibility incurred by artifacts in the image, clarity evaluates the relative visibility of veins and lesions in the image, and field definition evaluates the broad field of view of key retinal structures (e.g. the optic cup and disc, and fovea). Of note, three images categorized as ungradable by AutoMorph, were moved into our dataset on the basis of sufficient manual grading and viewing of the retinal vasculature. Following manual selection, a total of 123 usable Parkinson's disease images from 84 subjects were found to have met our criteria for inclusion. For each PD image, a healthy control (HC) with no history of PD was matched according to their age and gender to prevent covariate bias using the aforementioned image quality selection guidelines. All other fundus images and corresponding subjects were excluded. This constitutes our binary-labeled *overall* dataset of 246 fundus images (123 PD, 123 HC) from 168 subjects (84 PD, 84 HC). Lastly, we form two subsets of the data corresponding to a *prevalent dataset* of 154 fundus images (77 PD, 77 HC) from 110 subjects (55 PD, 55 HC) and an *incident* dataset of 92 fundus images (46 PD, 46 HC) from 58 subjects (29 PD, 29 HC).

Data pre-processing and model training

This study explores both deep learning and conventional machine learning models. Specifically, our deep learning models are convolutional neural networks (CNN) including AlexNet⁴⁰, VGG-16⁴¹, GoogleNet⁴², Inception-V3⁴³, and ResNet-50⁴⁴. Our conventional machine learning models include Logistic Regression, Elastic Net (regularized with hyperparameters $\lambda_{L_1} = 1$, $\lambda_{L_2} = 0.5$, where L_1 and L_2 designate the respective norms), and support vector machines (linear and radial basis function kernels).

Each of our models is purposed for binary classification with ground-truth labels of 0 (HC) or 1 (PD). A square crop preprocessing around the fundus image was employed to remove the external border. All fundus images are set to the left-eye view orientation to remove bias in the spatial orientation. To increase the diversity

of our dataset, we apply spatial rotations, horizontal, and vertical augmentations with a probability of 0.4. Images are inputted as a tensor for deep learning models (e.g. an RGB volume) and as a flattened array for conventional machine learning models. For computational efficiency, images are resized to 256×256 with 3 RGB channels or 1 channel (for gray-scale conversion and vessel segmented images).

Our deep learning models are initialized with ImageNet classification weights, with the input data normalized to the mean and standard deviation of ImageNet. The input of our conventional machine learning models is standardized to unit variance according to the training set's mean and standard deviation. Our testing evaluation is designed using a five-fold stratified cross-validation, wherein each Parkinson's subject is exactly matched to its corresponding age and gender matched control to remove bias of age and gender within the training and evaluation of our models. To further emphasize our training procedure, one-fold is selected as a test set, while the remaining four folds are delegated for allowable training. Of the four folds, one is selected as an internal validation set for hyper-parameter tuning while the model is initially trained using three of the training folds. For our conventional machine learning models, the regularization C-parameter in the set [0.1, 0.01, 1, 10, 100] is optimized. On the other hand, we fine-tune our deep learning models pre-trained on ImageNet using a binary cross-entropy loss, Adam optimizer⁴⁵, learning rate of $1e-4$, batch size of 64, and 100 epochs with early stopping. The choice of the learning rate was chosen to be roughly optimal and stable upon early experimentation of learning rates ranging from $1e-1$ to $1e-4$. After the discovery of the optimal hyper-parameters, the model re-consolidates the training set as the collection of all four training folds, and evaluated on the final test-set.

Performance evaluation

The performance of our PD classifiers is averaged over five randomized repetitions of five-fold stratified cross-validation, for a total of 25 testing evaluations. We consider the area under the receiver operating characteristic curve (AUC), accuracy, positive predictive value (PPV), negative predictive value (NPV), sensitivity (true positive rate), specificity (true negative rate), and the F1-score. Due to the lack of longitudinal data for the same subject, the metrics are evaluated on a per-image basis, treating each sample independently without data leakage (subject-level cross validation-level split). For reference, the training time, testing time, and number of model parameters are recorded (Supplementary Table 4).

Explainability evaluation

Qualitative visual explanations for our deep learning model predictions are accomplished by the Guided Backpropagation⁴⁶ algorithm, allowing an interpretable heat map of significant features. Quantitative explanations are provided by the explanation infidelity (INFD) and explanation sensitivity (SENS) metrics (to be distinguished from the classification performance metric *sensitivity*), where a lower result for both INFD and SENS provides evidence for better-robust models. The explanation infidelity and explanation sensitivity are computed at test-time over 50 perturbations drawn from a probability distribution $N \sim (0, 0.01^2)$, the definitions of which are shown below, and the details discussed in the corresponding paper⁴⁷.

$$\text{SENS}_{\text{MAX}}(\phi, \mathbf{f}, \mathbf{x}, r) = \max_{\|\mathbf{y}-\mathbf{x}\| \leq r} \|\phi(\mathbf{f}, \mathbf{y}) - \phi(\mathbf{f}, \mathbf{x})\|$$

$$\text{INFD}(\phi, \mathbf{f}, \mathbf{x}) = \mathbb{E}_{\mathbf{I}} \left[\left(\mathbf{I}^T \phi(\mathbf{f}, \mathbf{x}) - (\mathbf{f}(\mathbf{x}) - \mathbf{f}(\mathbf{x} - \mathbf{I})) \right)^2 \right]$$

where ϕ is an explanation function (e.g., Guided Backpropagation), \mathbf{f} is a black-box model (e.g., CNN), \mathbf{x} is in the input (e.g., fundus image), r is the input neighborhood radius, and \mathbf{I} is the perturbation, here drawn from the noise distribution $N \sim (0, 0.01^2)$.

Data availability

This research has been conducted using the UK Biobank, a publicly accessible database, under Application Number 48388. The datasets are available to researchers through an open application via <https://www.ukbiobank.ac.uk/register-apply/>.

Code availability

The underlying code for this study may be given permission by the authors from the GitHub repository <https://github.com/lab-smile/RetinaPD> but permission to the data is restricted to UKB applicants.

Received: 2 May 2023; Accepted: 10 February 2024

Published online: 13 February 2024

References

1. Yang, W. *et al.* Current and projected future economic burden of Parkinson's disease in the US. *NPJ Parkinson's Dis.* **6**, 15 (2020).
2. World Health Organization. Parkinson disease. <https://www.who.int/news-room/fact-sheets/detail/parkinson-disease> (2022).
3. Tolosa, E., Garrido, A., Scholz, S. W. & Poewe, W. Challenges in the diagnosis of Parkinson's disease. *Lancet Neurol.* **20**, 385–397 (2021).
4. Damier, P., Hirsch, E. C., Agid, Y. & Graybiel, A. M. The substantia nigra of the human brain. *Brain* **122**, 1437–1448 (1999).
5. Meissner, W. G. When does Parkinson's disease begin? From prodromal disease to motor signs. *Revue Neurologique* **168**, 809–814 (2012).
6. Postuma, R. B. *et al.* MDS clinical diagnostic criteria for Parkinson's disease: MDS-PD clinical diagnostic criteria. *Mov. Disord.* **30**, 1591–1601 (2015).

7. Martin, W. R. W. *et al.* Is levodopa response a valid indicator of Parkinson's disease?. *Mov. Disord.* **36**, 948–954 (2021).
8. Jankovic, J. Motor fluctuations and dyskinesias in Parkinson's disease: Clinical manifestations. *Mov. Disord.* **20**, S11–S16 (2005).
9. De La Fuente-Fernández, R. Role of DaTSCAN and clinical diagnosis in Parkinson disease. *Neurology* **78**, 696–701 (2012).
10. Postuma, R. B. *et al.* Validation of the MDS clinical diagnostic criteria for Parkinson's disease. *Mov. Disord.* **33**, 1601–1608 (2018).
11. Beach, T. G. & Adler, C. H. Importance of low diagnostic accuracy for early Parkinson's disease. *Mov. Disord.* **33**, 1551–1554 (2018).
12. London, A., Benhar, I. & Schwartz, M. The retina as a window to the brain—from eye research to CNS disorders. *Nat. Rev. Neurol.* **9**, 44–53 (2013).
13. Gupta, S., Zivadinov, R., Ramanathan, M. & Weinstock-Guttman, B. Optical coherence tomography and neurodegeneration: Are eyes the windows to the brain?. *Expert Rev. Neurother.* **16**, 765–775 (2016).
14. Archibald, N. K., Clarke, M. P., Mosimann, U. P. & Burn, D. J. The retina in Parkinson's disease. *Brain* **132**, 1128–1145 (2009).
15. Mohana Devi, S., Mahalaxmi, I., Aswathy, N. P., Dhivya, V. & Balachandar, V. Does retina play a role in Parkinson's Disease?. *Acta Neurologica Belgica* **120**, 257–265 (2020).
16. Diederich, N. J., Raman, R., Leurgans, S. & Goetz, C. G. Progressive worsening of spatial and chromatic processing deficits in Parkinson disease. *Arch. Neur.* **59**, 1249 (2002).
17. Kwapong, W. R. *et al.* Retinal microvascular impairment in the early stages of Parkinson's disease. *Investig. Ophthalmol. Vis. Sci.* **59**, 4115 (2018).
18. Hajee, M. E. Inner retinal layer thinning in Parkinson disease. *Arch. Ophthalmol.* **127**, 737 (2009).
19. Mailankody, P. *et al.* Optical coherence tomography as a tool to evaluate retinal changes in Parkinson's disease. *Parkinsonism Relat. Disord.* **21**, 1164–1169 (2015).
20. Kromer, R. *et al.* Evaluation of retinal vessel morphology in patients with Parkinson's disease using optical coherence tomography. *PLoS ONE* **11**, e0161136 (2016).
21. Dai, L. *et al.* A deep learning system for detecting diabetic retinopathy across the disease spectrum. *Nat. Commun.* **12**, 3242 (2021).
22. Yan, Q. *et al.* Deep-learning-based prediction of late age-related macular degeneration progression. *Nat. Mach. Intell.* **2**, 141–150 (2020).
23. Sudlow, C. *et al.* UK Biobank: An open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS Med.* **12**, e1001779 (2015).
24. Terracciano, A. *et al.* Neuroticism and risk of Parkinson's disease: A meta-analysis. *Mov. Disord.* **36**, 1863–1870 (2021).
25. Xu, Q. *et al.* Diabetes and risk of Parkinson's disease. *Diabetes Care* **34**, 910–915 (2011).
26. Kizza, J. *et al.* Cardiovascular risk factors and Parkinson's disease in 500,000 Chinese adults. *Ann. Clin. Transl. Neurol.* **6**, 624–632 (2019).
27. Zhou, Y. *et al.* AutoMorph: Automated retinal vascular morphology quantification via a deep learning pipeline. *Trans. Vis. Sci. Tech.* **11**, 12 (2022).
28. Carelli, V., La Morgia, C., Ross-Cisneros, F. N. & Sadun, A. A. Optic neuropathies: The tip of the neurodegeneration iceberg. *Hum. Mol. Genet.* **26**, R139–R150 (2017).
29. Pilat, A. *et al.* In vivo morphology of the optic nerve and retina in patients with Parkinson's disease. *Investig. Ophthalmol. Vis. Sci.* **57**, 4420 (2016).
30. Miri, S. *et al.* The avascular zone and neuronal remodeling of the fovea in Parkinson disease. *Ann. Clin. Transl. Neurol.* **2**, 196–201 (2015).
31. Ascherio, A. & Schwarzschild, M. A. The epidemiology of Parkinson's disease: Risk factors and prevention. *Lancet Neurol.* **15**, 1257–1272 (2016).
32. Miller, I. N. & Cronin-Golomb, A. Gender differences in Parkinson's disease: Clinical characteristics and cognition. *Mov. Disord.* **25**, 2695–2703 (2010).
33. Hu, W. *et al.* Retinal age gap as a predictive biomarker of future risk of Parkinson's disease. *Age Ageing* **51**, afac062 (2022).
34. Nunes, A. *et al.* Retinal texture biomarkers may help to discriminate between Alzheimer's, Parkinson's, and healthy controls. *PLoS ONE* **14**, e0218826 (2019).
35. Tian, J. *et al.* Modular machine learning for Alzheimer's disease classification from retinal vasculature. *Sci. Rep.* **11**, 238 (2021).
36. Wisely, C. E. *et al.* Convolutional neural network to identify symptomatic Alzheimer's disease using multimodal retinal imaging. *Br. J. Ophthalmol.* **106**, 388–395 (2022).
37. Bush, K., Rannikmae, K., Wilkinson, T., Schnier, C. & Sudlow, C. Definitions of Parkinson's disease and the major causes of parkinsonism, UK Biobank Phase 1 Outcomes Adjudication. (2018).
38. Chua, S. Y. L. *et al.* Cohort profile: Design and methods in the eye and vision consortium of UK Biobank. *BMJ Open* **9**, e025077 (2019).
39. Fu, H. *et al.* Evaluation of Retinal Image Quality Assessment Networks in Different Color-Spaces. vol. 11764, 48–56 (2019).
40. Krizhevsky, A., Sutskever, I. & Hinton, G. E. ImageNet classification with deep convolutional neural networks. *Commun. ACM* **60**, 84–90 (2017).
41. Simonyan, K. & Zisserman, A. Very Deep Convolutional Networks for Large-Scale Image Recognition. Preprint at <http://arxiv.org/abs/1409.1556> (2015).
42. Szegedy, C. *et al.* Going Deeper with Convolutions. Preprint at <http://arxiv.org/abs/1409.4842> (2014).
43. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J. & Wojna, Z. Rethinking the Inception Architecture for Computer Vision. Preprint at <http://arxiv.org/abs/1512.00567> (2015).
44. He, K., Zhang, X., Ren, S. & Sun, J. Deep Residual Learning for Image Recognition. Preprint at <http://arxiv.org/abs/1512.03385> (2015).
45. Kingma, D. P. & Ba, J. Adam: A Method for Stochastic Optimization. Preprint at <http://arxiv.org/abs/1412.6980> (2017).
46. Selvaraju, R. R. *et al.* Grad-CAM: Visual explanations from deep networks via gradient-based localization. *Int J Comput Vis* **128**, 336–359 (2020).
47. Yeh, C.-K., Hsieh, C.-Y., Suggala, A. S., Inouye, D. I. & Ravikumar, P. On the (In)Fidelity and Sensitivity for Explanations. Preprint at <http://arxiv.org/abs/1901.09392> (2019).

Acknowledgements

This study was supported by NSF IIS 2123809.

Author contributions

R.F. and C.T. had full access to all the data in the study and take responsibility for the integrity of the data and the accuracy of the data analysis. Study concept and design: C.T., K.S., R.F. Acquisition, analysis, or interpretation of data: C.T., K.S., K.L., R.F. Drafting of the manuscript: C.T., K.S., K.L., Y.L., R.F. Critical revision of the manuscript for important intellectual content: C.T., R.F. A.R.-Z., J.C., A.A. Statistical analysis: C.T., K.S., Y.L. Obtained funding: R.F. Administrative, technical, or material support: R.F. Study supervision: R.F.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-024-54251-1>.

Correspondence and requests for materials should be addressed to R.F.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024