



A novel fusion architecture for detecting Parkinson's Disease using semi-supervised speech embeddings



Tariq Adnan^{1,3}, Abdelrahman Abdelkader^{1,3}, Zipei Liu¹, Ekram Hossain¹, Sooyong Park¹,
Md Saiful Islam^{1,4} & Ehsan Hoque^{1,2,4} ✉

We introduce a framework for screening Parkinson's disease (PD) using English pangram utterances. Our dataset includes 1306 participants (392 with PD) from both home and clinical settings, covering diverse demographics (53.2% female). We used deep learning embeddings from Wav2Vec 2.0, WavLM, and ImageBind to capture speech dynamics indicative of PD. Our novel fusion model for PD classification aligns different speech embeddings into a cohesive feature space, outperforming baseline alternatives. In a stratified randomized split, the model achieved an AUROC of 88.9% and an accuracy of 85.7%. Statistical bias analysis showed equitable performance across sex, ethnicity, and age subgroups, with robustness across various disease durations and PD stages. Detailed error analysis revealed higher misclassification rates in specific age ranges for males and females, aligning with clinical insights. External testing yielded AUROCs of 82.1% and 78.4% on two clinical datasets, and an AUROC of 77.4% on an unseen general spontaneous English speech dataset, demonstrating versatility in natural speech analysis and potential for global accessibility and health equity.

The diagnosis of Parkinson's Disease (PD) is traditionally reliant on the clinical assessments focused on the motor symptoms of the individuals¹. Traditional methods, while effective, often miss the subtle early symptoms of the disease, leading to delayed interventions². The situation is further exacerbated by the limited accessibility to specialized neurological health-care, particularly in regions with significantly lower ratios of neurologists to the population. For instance, Bangladesh had only 86 neurologists for over 140 million people in 2014³, while some African nations had one neurologist per three million people, with 21 countries having fewer than five neurologists each⁴. Given the expected doubling of PD cases by 2030⁵, there's a pressing need for accessible, home-based diagnostic solutions to address global disparities in healthcare access.

Recent advancements have seen a shift towards integrating digital biomarkers to develop automated AI based at-home PD detection and progression tracking tools^{6–10}. Techniques vary from sensor-based nocturnal breathing signal⁹ and accelerometric data collection¹¹, to digital analysis of facial expressions¹². However, wearables and sensors may be inconvenient for the elderly^{13,14}, and posed expressions can miss subtle diagnostic cues.

Alternatively, speech analysis offers a non-invasive route, leveraging natural speech patterns for PD detection. Traditional speech analysis in PD—primarily relied on sustained phonation tasks^{6,15–21}—

although useful, does not reflect the complexities of natural speech. To counter that, studies have been proposed to use continuous speech to develop PD classifiers using varying technologies such as CNNs²², time-frequency analysis²³, or SVMs²⁴. However, these studies rely on fixed recording setups and small sample sizes, which limits the generalizability of the models and fail to adequately address accessibility concerns. Even with larger datasets, models like CNNs and SVMs face structural limitations. CNNs, while powerful for feature extraction, are primarily designed for spatial data, and their application to time-series or speech data can be challenging unless properly adapted^{25,26}. They may require deep architectures to capture complex temporal dependencies in PD speech data, increasing the risk of overfitting if the model is not properly regularized or if the dataset lacks sufficient variability to cover real-world scenarios²⁷. SVMs, although effective on small, well-separable data, may struggle in high-dimensional, noisy feature space²⁸ common in speech signals²⁹ due to their limitation in temporal pattern handling compared to neural networks³⁰. These structural lacking of the modeling architectures in the existing literature underscore the demand of more adaptable models, which can capture the nuanced vocal changes associated with PD while generalizing better across diverse recording conditions and larger datasets.

¹Department of Computer Science, University of Rochester, Rochester, NY, USA. ²Ministry of Defense Health Services, Riyadh, Saudi Arabia.

³These authors contributed equally: Tariq Adnan, Abdelrahman Abdelkader. ⁴These authors jointly supervised this work: Md Saiful Islam, Ehsan Hoque. ✉e-mail: mehoque@cs.rochester.edu



Traditional feature engineering, while interpretable, requires extensive time, human effort, and domain expertise to identify and select relevant features manually³¹. This approach introduces the risk of human bias, as researchers may prioritize familiar or well-documented features, potentially overlooking subtle yet critical nuances present in neurological speech patterns. Studies have shown that even with careful feature selection, critical variations indicating early-stage PD can be missed, reinforcing the need for more automated, data-driven feature extraction techniques, such as those used in deep learning models^{32–35}. Recent advancements in semi-supervised learning models, such as WavLM³⁶, Wav2Vec 2.0³⁷, and ImageBind³⁸, offer pre-trained, openly accessible models that capture complex and abstract representations of speech data. These models, trained on diverse, large-scale datasets, exhibit strong performance in several downstream applications like automatic speech recognition (ASR)^{37,39} and speech diarization³⁶, suggesting their embeddings capture intricate, high-dimensional speech dynamics. These embeddings hold significant promise to reveal subtle voice changes associated with the PD that handcrafted features might miss. Despite their potential, only limited works⁴⁰ have been done to explore semi-supervised acoustic models' utility in PD classification. Although this study⁴⁰ was conducted only on a small cohort of 60 participants (28 PD), it found significant promise in applying semi-supervised speech embeddings from wav2vec in PD classification for both English and Italian language. On the other hand, fusion after projection of one feature set into another latent space can achieve significant improvement in feature alignment, noise reduction, and dimensionality consistency^{13,41,42}, simplifying the architecture and enhancing cross-modal learning capacity^{43–45}. Despite the significant advancements in speech processing and disease classification, the literature has yet to fully explore the potential of combining semi-supervised deep speech embeddings with projection-based fusion models for PD classification.

In our study, we leveraged a large-scale dataset of 1306 participants, including 392 PD-diagnosed individuals, collected from varied environments—participants' homes, clinical settings, and PD care centers—enabling us to

overcome the limitations of smaller cohorts and constrained recording environments in PD detection. We expanded the application of semi-supervised vector embeddings of free-flow speech through a novel projection-based fusion approach, demonstrating that fusing deep embeddings from WavLM to ImageBind feature space enhances the model's ability to capture PD-specific speech characteristics. This fusion approach notably outperformed traditional concatenation by effectively aligning embeddings, reducing redundancy, and enabling a more nuanced representation of disease indicators within speech. To prepare the dataset, participants were instructed to record themselves while uttering the English pangram that starts with “the quick brown fox.” We modeled the free-flow speech using the pangram utterance to simplify the analysis as it ensures a consistent speech content across all participants. This study highlights that semi-supervised embeddings are highly effective for PD classification when paired with projection-based fusion, marking a key contribution in identifying PD-specific vocal nuances often missed by handcrafted features. We maximized the generalizability and robustness of the model by collecting data across diverse environments—homes, clinics, and a PD care facility—and a demographically diverse cohort. We conducted extensive error analysis using multiple methodologies to identify any comparatively underperforming demographic subgroups. With our web-based framework, English-speaking individuals having access to a webcam-enabled laptop or desktop can record their speech and receive a PD screening. This approach can enhance accessibility, particularly in regions with limited access to neurological care. Figure 1 illustrates our PD classification pipeline, which processes raw video inputs to determine PD presence effectively.

Results

Dataset

We collected our dataset from 1306 participants, comprising 392 of them diagnosed with PD and 914 without the condition. We used PARK⁴⁶, a web-based framework for data collection purpose. Using the framework, each participant recorded themselves in front of the web camera while reciting

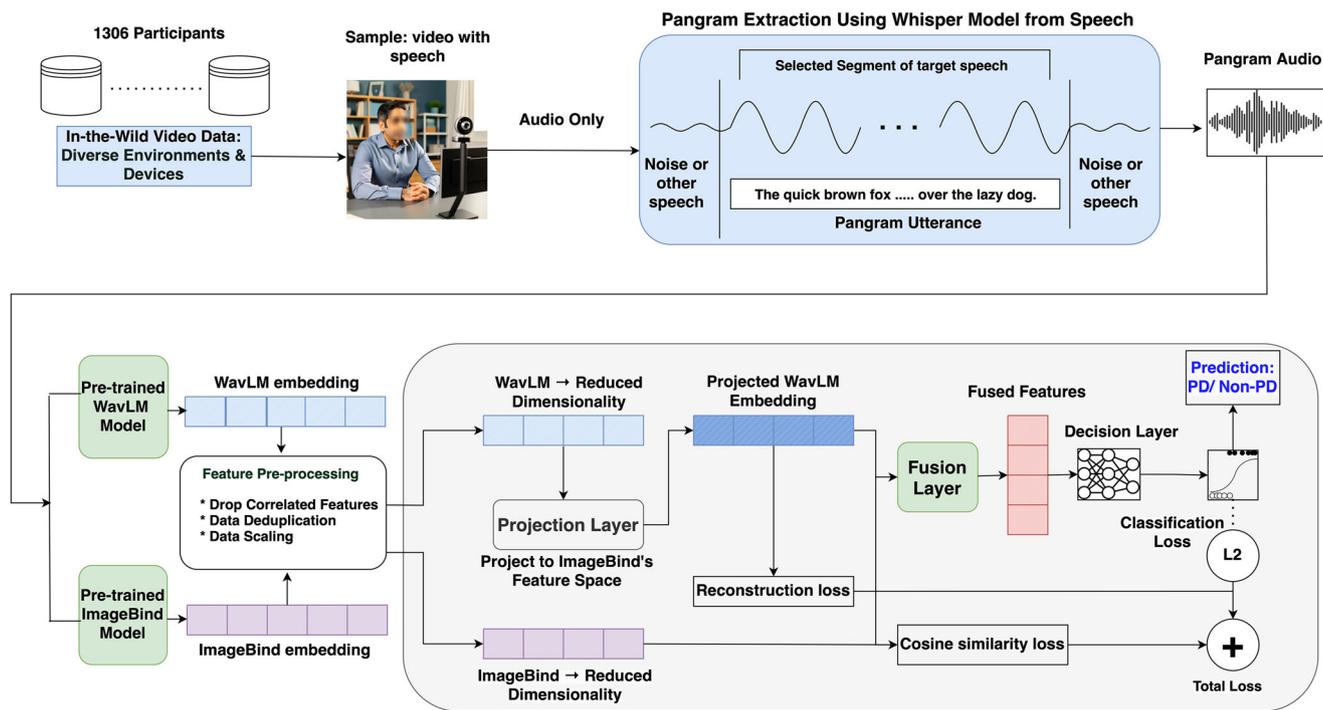


Fig. 1 | Our proposed framework of fusion based PD classifier using deep embeddings from WavLM and ImageBind. First, the speech is separated from video datasets. Then the segment of the audio file where the participants utter the pangram is separated. Vector embeddings from the last layers of WavLM and

ImageBind are extracted for the speech data. Then WavLM features are projected into the space of ImageBind features set. Finally the projected features are fused and passed through a classification layer that can determine the participant as PD or control. Note that the image of the person is AI generated.

Table 1 | Demographic information of the participants

Demographic Properties	Attribute Number of Participants	With PD 392 (30.0%)	Without PD 914 (70.0%)	Total 1306 (100.0%)
Sex	Female	171 (43.6%)	524 (57.3 %)	695(53.2%)
	Male	218 (55.6%)	390 (42.7 %)	608(46.6%)
	Nonbinary	1 (0.3%)	0 (0 %)	1(0.1%)
	Unknown	2 (0.5%)	0 (0 %)	2(0.2%)
Age in years range: 16.0–93.0 mean: 62.2, std: 12.7	Below 20	0 (0.0%)	8 (0.9%)	8 (0.61%)
	20–29	2 (0.5%)	41 (4.5%)	43 (3.29%)
	30–39	3 (0.8%)	24 (2.6%)	27 (2.07%)
	40–49	9 (2.3%)	20 (2.2%)	29 (2.22%)
	50–59	40 (10.2%)	160 (17.5%)	200 (15.31%)
	60–69	122 (31.1%)	349 (38.2%)	471 (36.06%)
	70–79	142 (36.2%)	107 (11.7%)	249 (19.07%)
	80 and above Unknown	23 (5.9%) 51 (13.0%)	7 (0.8%) 198 (21.7%)	30 (2.3%) 249 (19.07%)
Ethnicity	White	223 (56.9%)	638 (69.8%)	861 (65.9%)
	Black or African American	6 (1.5%)	42 (4.6%)	48 (3.7%)
	American Indian or Alaska Native	1 (0.3%)	4 (0.4%)	5 (0.4%)
	Asian	4 (1.0%)	55 (6.0%)	59 (4.5%)
	Others	2 (0.5%)	5 (0.5%)	7 (0.5%)
	Unknown	156 (39.8%)	170 (18.6%)	326 (25.0%)
Disease duration in years range: 1–27 mean: 6.32, std: 5.24	< =2	35(8.9%)	–	–
	2–5	42(10.7%)	–	–
	5–10	36(9.2%)	–	–
	10–15	19(4.8%)	–	–
	15–20	8(2.0%)	–	–
	> 20 Unknown	3(0.8%) 249(63.5%)	– –	– –
Hoehn and Yahr PD stages	0: Asymptomatic	65(16.6%)	–	–
	1: Unilateral involvement only	41(10.5%)	–	–
	2: Bilateral involvement, no balance issues	57(14.5%)	–	–
	3: Moderate symptoms, postural instability	50(12.8%)	–	–
	Unknown	179(45.7%)	–	–
Recording Environment	Home Recorded	67 (17.1%)	585 (64.0%)	652 (49.9%)
	Clinical Setup	117 (29.8%)	235 (25.7%)	352 (27.0%)
	PD Care Facility	185 (47.2%)	85 (9.3%)	270 (20.7%)

the “quick brown fox” English pangram across three distinct recording environments—Home Recorded, Clinical Setup, and PD Care Facility—each varying in terms of ambient noise and data collection setting and/or equipment. At the Clinical Setup, and PD Care Facility, some participants provided multiple data samples and eventually, we obtained 1, 854 video clips having audio of pangram utterance. Note that PD labels of the participants from Clinical Setup and PD Care Facility cohorts are clinically validated, while the Home Recorded labels are self-reported. The demographic information of the participating cohort is detailed in Table 1.

Feature Extraction

In this study, we aim to capture the nuanced speech dynamics indicative of PD using advanced deep learning embeddings from pangram utterance speech. We utilized three state-of-the-art semi-supervised speech models: Wav2Vec 2.0³⁷, WavLM³⁶, and ImageBind³⁸ to extract intermediate vector representations from their last hidden layers, capturing sophisticated and informative features of the speech data. Alongside, to assess the efficacy of these deep embedding features relative to traditional acoustic features, we

extracted classical features using the methodology proposed by Rahman et al.⁶. For comprehensive analysis, we compiled four distinct feature sets from our audio datasets: 39-dimensional acoustic features, 768-dimensional Wav2Vec 2.0 features, 1024-dimensional WavLM features, and 1024-dimensional ImageBind features. We then trained various deep learning models using these feature sets to distinguish between individuals with and without PD, exploring how different types of features contribute to the model’s performance.

Performance Evaluation on Standard Train-Validation-Test Split

In our initial experiment, we combined all three data cohorts, segmenting the dataset into three random splits: 70% for training, and 15% each for validation and testing. The validation set was used to select the best-performing model for subsequent testing. Table 2(a) provides the data split details and demographic distribution across the subsets, where we can observe that each subset has a fairly balanced representation of demographic subgroups. Importantly, the split was conducted based on participants’ IDs, ensuring that all data samples from each participant were grouped into

Table 2 | Dataset Demographics for Model Evaluation Splits

(a)							
Cohort	# of Participants	# of PD	# of Female	# of Non-White	# of < 50	# of Clinical	# of PD Care
Training Set	916 (70.1%)	284 (31.0%)	481 (52.5%)	82 (9.0%)	71 (7.8%)	291 (31.8%)	214 (23.4%)
Validation Set	195 (14.9%)	54 (27.7%)	116 (59.5%)	19 (9.7%)	17 (8.7%)	25 (12.8%)	30 (15.4%)
Test Set	195 (14.9%)	54 (27.7%)	98 (50.3%)	18 (9.2%)	19 (9.7%)	36 (18.5%)	26 (13.3%)
(b)							
Cohort	# of Participants	# of PD	# of Female	# of Non-White	# of < 50		
Training Set: Home Recorded and PD Care Facility	771 (59.0%)	225 (29.2%)	432 (56.0%)	94 (12.2%)	77 (10.0%)		
Validation Set: Home Recorded and PD Care Facility	183 (14.0%)	50 (27.3%)	105 (57.4%)	19 (10.4%)	21 (11.5%)		
Test Set: Clinical Setup	352 (27.0%)	117 (33.2%)	158 (44.9%)	3 (0.9%)	9 (2.6%)		
Training Set: Home Recorded and Clinical Setup	837 (64.1%)	171 (20.4%)	448 (53.5%)	83 (9.9%)	67 (8.0%)		
Validation Set: Home Recorded and Clinical Setup	199 (15.2%)	36 (18.1%)	106 (53.3%)	20 (10.1%)	15 (7.5%)		
Test Set: PD Care Facility	270 (20.7%)	185 (68.5%)	141 (52.2%)	13 (4.8%)	25 (9.3%)		

This table presents the demographic breakdown of participants across different splits for model evaluation: (a) using a conventional random train-validation-test split, and (b) using a cross-environment split where Clinical Setup and PD Care Facility recording environments alternately serve as the test sets, with the remaining environments combined for training and validation.

either the train, validation, or test sets. Furthermore, the stratified splitting method guaranteed fair representation of PD participants across all three sets.

First, we developed multiple baseline models to benchmark our approach, incorporating both traditional and deep embedding features. First, we implemented a CNN-based model trained directly on raw speech data as an end-to-end PD classification pipeline. This model, however, underperformed with an AUROC of 59.18% and accuracy of 63.71%, highlighting the challenges in capturing PD-specific speech characteristics directly from time-series audio data. CNNs are architecturally constrained by their local receptive fields, which limits their ability to capture both global context and long-range dependencies in the input data^{47,48}, leading to sub-optimal performance in this task. Additionally, we tested both support vector machine (SVM) and neural network classifiers across four distinct feature sets: classical acoustic features, WavLM, Wav2Vec2, and ImageBind embeddings. Among the SVM-based models using four different feature sets, the one trained on WavLM embeddings demonstrated the best performance, achieving an AUROC of 75.43% and accuracy of 69.65%, surpassing the SVM model with classical acoustic features, which reached an AUROC of 74.82% and accuracy of 67.94%.

Our neural network-based classifier, particularly when trained with WavLM embeddings, significantly outperformed all other baseline models. It achieved an AUROC of 85.89% and an accuracy of 81.01%, surpassing the closest competitor—model trained with ImageBind features—by 5% in AUROC. Despite a comparatively lower sensitivity of 56.25%, the model excelled in other metrics, achieving a specificity of 90.63%, a Positive Predictive Value (PPV) of 81.01%, and the Negative Predictive Value (NPV) of 80.79%. Notably, all models trained with deep embedding features consistently surpassed those using traditional acoustic features, and the neural network-based classifiers outperformed the SVM models in general. Please refer to the first segment of six rows in Table 2 for a summary of the evaluation metrics across various baseline models.

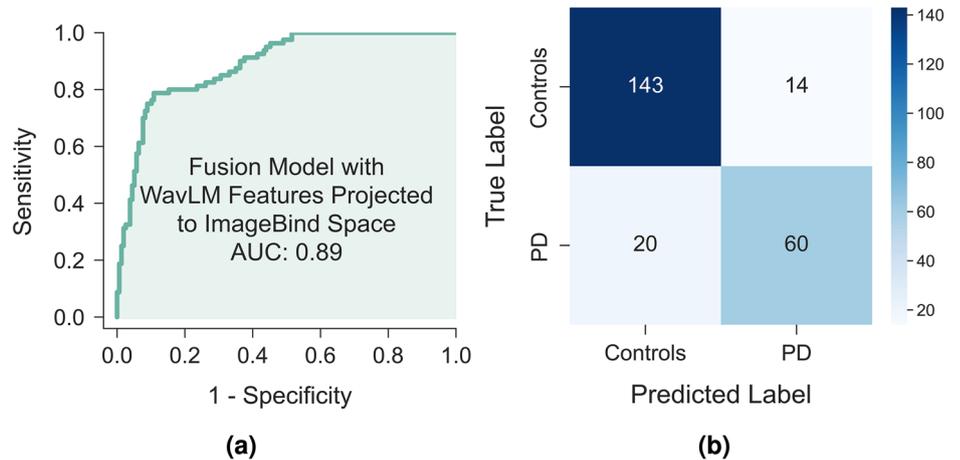
To enhance model performance by leveraging the complementary strengths of different feature sets, we developed fusion models using concatenation and projection-based approaches. By concatenating four different feature sets in all possible combinations—resulting in 11 unique set—we observed consistent improvements, especially with combinations involving WavLM features. The best results were achieved when combining all three deep embeddings (Wav2Vec2, WavLM, and ImageBind), with an AUROC of 89.49% and accuracy of 82.28%. Although specificity (85.99%) and PPV (73.17%) were slightly lower than the best baseline model, this

approach significantly improved sensitivity (from 56.25% to 75%) and NPV (from 80.79% to 87.10%).

In recent years, projection-based fusion has emerged as a powerful approach for enhancing representation learning in classifiers. This method involves projecting features from one feature space into the space of another feature set, thus optimizing the use of complementary information while minimizing redundancy. While our early fusion models using concatenation showed promising improvements, our projection-based fusion models further refined PD classification. Despite a slight decrease in AUROC (88.94% compared to the previous best of 89.49%), the fusion model projecting WavLM features into the feature space of ImageBind features achieved a significant increase in accuracy to 85.65% from 82.28%. Additionally, it outperformed all other models (or achieved the similar best performance), in terms of other evaluation metrics with a sensitivity of 75.00%, specificity of 91.08%, PPV of 81.08%, and NPV of 87.73%. Figure 2 demonstrates the ROC curve and confusion matrix of this best-performing fusion model. To further optimize model performance, we experimented with threshold tuning on the validation set and applied the selected threshold to the test set. We found that lowering the threshold from the default 0.50 to an optimal value of 0.44 provided a flexible trade-off between sensitivity and specificity, allowing for adjustments based on deployment requirements. While most key metrics remained consistent, this adjustment led to a slight increase in F1-score (from 77.92% to 78.75%) and an improvement in sensitivity (from 75.00% to 78.75%) at the cost of a minor decrease in specificity (from 91.08% to 89.17%). This trade-off suggests that, in practice, the model’s operating point can be fine-tuned based on the requirements of the deployment site—for instance, prioritizing higher sensitivity to detect more PD cases while accepting a slightly higher false positive rate. The full analysis, including threshold selection curves, is provided in the Supplementary Note 5.

We also investigated whether reducing the dimensionality of the deep embeddings generated by WavLM and ImageBind (1024 dimensions each) could lead to a more efficient model. While our default best model utilizes the full-dimensional WavLM and ImageBind embeddings and requires only 58.8 ± 4.2 milliseconds per epoch for training due to its simplistic architecture, we applied Principal Component Analysis (PCA)^{49,50} to explore the potential benefits of dimensionality reduction. With this approach, the WavLM features were reduced to 192 dimensions, and the ImageBind features to 280 dimensions, while retaining at least 98% of the original variance for both modalities. This dimensionality reduction resulted in a notable reduction in training time (41.6 ± 2.5 milliseconds per epoch).

Fig. 2 | Performance evaluation of PD classifiers from speech in a random split configuration. **a** and **b** respectively demonstrate the AUROC curve and the confusion matrix of our best performing novel fusion model, which projects WavLM features into the feature space of ImageBind features.



However, the performance of the reduced model was significantly lower across all major evaluation metrics, demonstrating that the reduction in training time was not worth the trade-off in performance. The PCA-reduced model achieved an accuracy of 75.95% and an AUROC of 82.22%. While the model maintained a high specificity (94.27%), its sensitivity dropped to just 40.00%. Since PCA primarily focuses on finding directions of maximum variance in the data, it may not preserve the features most relevant for distinguishing between PD and Non-PD, leading to a significant drop in overall model performance.

Since our dataset was imbalanced, we evaluated our best fusion model on a resampled dataset using three different techniques: SMOTE⁵¹, Random Undersampling⁵², and Random Oversampling⁵². Among these, SMOTE yielded the best performance. However, even with the balanced dataset generated by synthetic samples, SMOTE failed to outperform the model developed without addressing data imbalance, achieving only an AUROC of 74.45%.

Additionally, beyond evaluating the model on a randomly selected test set, we performed a 10-fold cross-validation on our best-performing (on the randomly selected test set) model configuration. This configuration yielded similar performance, with an AUROC of 90.86% and an accuracy of 85.37%. It also demonstrated 81.88% sensitivity, 89.53% specificity, 80.29% PPV, and 88.63% NPV, further reinforcing the model’s robustness.

Finally, ImageBind inherently supports multimodal input and we had collected video data for the speech task. Although our key analyses are focused on audio modality in order to consistently evaluate the state-of-the-art semi-supervised speech representation models, we conducted preliminary experiments to explore the potential of using multimodal (audio + video) embeddings from the ImageBind model. As an initial exploration, we replaced ImageBind’s audio-only embeddings with multimodal embeddings in our baseline model, which notably outperformed the audio-only embeddings across all metrics. Specifically, it achieved an AUROC of 84.88% (4.46% improvement over ImageBind audio-only embeddings) and an accuracy of 78.48% (4.22% improvement), suggesting that incorporating video information may enhance performance. Motivated by this, we further explored two variants of our proposed fusion model: one where multimodal ImageBind embeddings are projected to WavLM before fusion and another where WavLM embeddings are projected to multimodal ImageBind embeddings before fusion. While these models demonstrated competitive performance, neither surpassed the best-performing speech-only model. A likely explanation is that multimodal ImageBind is less compatible with the WavLM embeddings, as their inputs are of different modalities. Therefore, although multimodal ImageBind embedding is better than audio-only ImageBind when experimented standalone, audio-only ImageBind embeddings yield better performance when fused with WavLM. However, these results highlight the potential promise of multimodal analysis for PD detection using the existing dataset.

The middle segment with ten rows in Table 3 summarizes the performance achieved by some of our top-performing fusion models.

Generalizability Test on External Datasets

To evaluate the model’s generalizability and its performance on datasets with probable distribution shift, we tested our best-performing fusion model using the datasets from Clinical Setup and PD Care Facility cohorts, while the training was conducted on the remaining cohorts, including the Home Recorded cohort. Due to a significant imbalance between PD and control participants in the Home Recorded cohort, with only 68 PD participants compared to 585 controls, it was not used in external testing. The demographic distribution of the training, validation, and test sets for this generalizability test is detailed in Table 2(b). In the first configuration, using the Home Recorded and Clinical Setup cohorts for training and testing on the PD Care Facility cohort, the model achieved an AUROC of 82.12% and accuracy of 74.69%. This represents a decline of 6.82% in AUROC and 10.96% in accuracy from the best random split performance. Additionally, the model showed a sensitivity of 71.61%, specificity of 82.42%, PPV of 91.11%, and a notably lower NPV of 53.58%.

In the second configuration, testing on the Clinical Setup cohort after training on the Home Recorded and PD Care Facility cohorts, the model recorded an AUROC of 78.43% and accuracy of 70.19%, down by 10.51% and 15.46% respectively. Sensitivity was 77.32%, specificity was 65.84%, PPV was 58.10%, and NPV remained comparable at 82.58%. Figure 3 demonstrate the ROC curve and confusion matrix when our best fusion model was tested on data from the Clinical Setup and PD Care Facility cohorts, respectively. The detailed evaluation metrics for these external datasets are shown in the last two rows of Table 3.

Error Analysis

To determine if the model is under performing for any particular demographic subgroup in the random split configuration, we performed rigorous error analysis in terms of four key demographic properties of the participants that were not included in the model’s training and validation stages: sex, ethnicity, age, and recording environment.

First, we performed statistical significance tests to evaluate whether the model’s performance differed significantly across complementary demographic subgroups. All statistical tests were conducted at an $\alpha = 0.05$ significance level. Given that we conducted six different hypothesis tests on the same evaluation set, we applied a Bonferroni correction⁵³ to control for multiple comparisons, adjusting the effective significance level to be $\alpha^* = \frac{\alpha}{6} = 0.0083$. First, we divided the dataset into multiple subgroups based on sex (Male vs. Female), ethnicity (White vs. Non-White), age (Below 50 years vs. 50 years and Above), and recording environment (Home Recorded vs. Clinical Setup vs. PD Care Facility). Participants missing demographic data were excluded from respective analyses. Note that in the

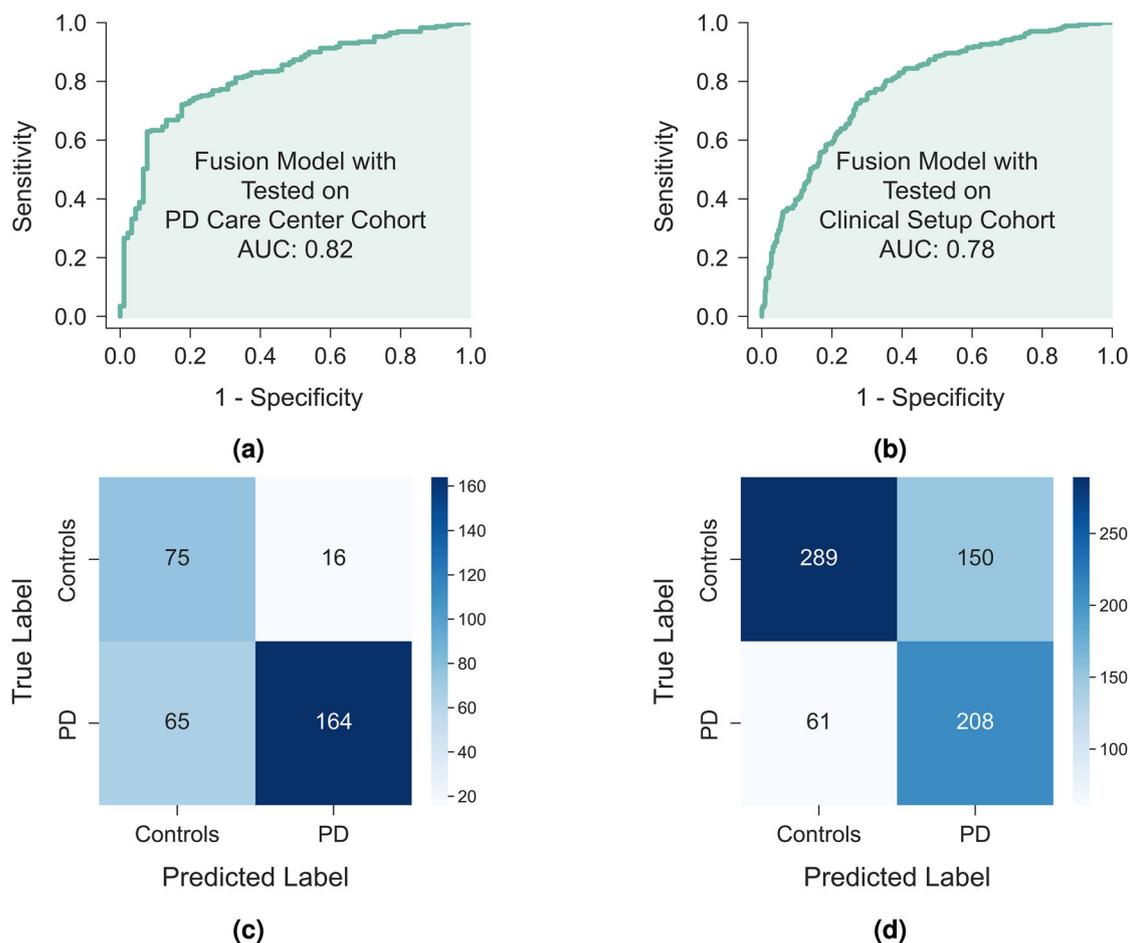


Fig. 3 | Performance evaluation of PD classifiers from speech on external test sets. **a** and **c** respectively demonstrate the AUROC curve and the confusion matrix of our best performing novel fusion model when tested on the dataset collected from PD

Care Facility. In contrast, **b** and **d** give such visualizations when the model was tested on the participating cohort from Clinical Setup.

random split experimental setting, our test set had 237 audio samples. To carry out the statistical tests, we employed Fisher’s Exact Test⁵⁴ to compare proportions between each pair of subgroups. This test was chosen due to its strength in handling categorical data and its suitability for our sample size, without the need to assume normality. Fisher’s Exact Test provides an exact *p*-value, ensuring the robustness of our results even in cases where the central limit theorem’s assumptions⁵⁵ for other parametric tests, like the *Z*-test, may not fully hold.

For the analysis based on sex, the model achieved an accuracy of 84.3% for the male subgroup (115 samples) and 86.9% for the female subgroup (122 samples). The *p*-value for this comparison was 0.5847, indicating that the difference in model performance between male and female participant groups was not statistically significant. Ethnically, the model achieved an accuracy of 82.5% for the White subgroup (183 samples) and 100.0% for the Non-White subgroup (18 samples). The *p*-value for this comparison was 0.0834, which was also not statistically significant. Age wise, the model achieved an accuracy of 100.0% for participants below 50 years old (19 samples) and 85.9% for participants aged 50 years and above (199 samples). With a *p*-value of 0.1423, again this difference was not statistically significant. Since we had three distinct recording environments, we conducted statistical test for each of the pairs. The model achieved an accuracy of 91.0% for the Home Recorded cohort (133 samples) and 72.2% for the Clinical Setup cohort (72 samples). The *p*-value for this comparison was 0.0010, indicating that the difference was statistically significant at the 95% level. For the comparison between Home Recorded (133 samples, 91.0% accuracy) and PD Care Facility cohorts (32 samples, 93.8% accuracy),

the *p*-value was 0.9211, showing no significant difference. Lastly, in the comparison between the Clinical Setup (72 samples, 72.2% accuracy) and PD Care Facility cohorts (32 samples, 93.8% accuracy) shows a *p*-value of 0.0176, which also came out to be statistically insignificant considering our corrected effective significance level. The outcome of different statistical significance test are summarized in Table 4.

Apart from the sub-group based statistical bias analysis, we also examined the influence of disease duration (in years) and Hoehn and Yahr (H&Y) PD stages on the predictive performance. First, we examined the association of disease duration and the model accuracy by using Spearman’s rank correlation among 110 samples with known durations. The correlation coefficient ($\rho = 0.18$) suggested a weak positive correlation between model accuracy and disease duration, which was not statistically significant ($p = 0.0579$). This indicates that the duration of the disease does not significantly affect the model’s effectiveness. Next, we analyzed the 75 audio samples from 35 participants in the test set with known PD stages and conducted a Kruskal-Wallis H test and Fisher’s exact test with Monte Carlo simulation (10,000 replications) to examine the association between PD stage categories and model accuracy. The PD stages among these 35 participants ranged from stage 0 (Asymptomatic) to stage 3 (Moderate symptoms, postural instability). The Kruskal-Wallis H test yielded a test statistic of 0.4595 with a *p*-value of 0.9276, while Fisher’s exact test resulted in a *p*-value of 0.9711. Both statistical analyses indicate no significant association between PD stage categories and the model’s predictive performance. See Supplementary Fig. 1 for a visual representation of model accuracy across different H&Y stages.

Table 3 | Performance reporting for the ablation studies in a summarized version

Experimental Setup	AUROC	Accuracy	Sensitivity	Specificity	PPV	NPV	F1 Score
SVM w/ Acoustic Features	74.82	67.94	64.19	69.93	53.06	78.67	58.10
SVM w/ WavLM Embeddings	75.43	69.65	64.19	72.50	55.31	79.28	59.42
CNN w/ raw speech	59.18	63.71	5.40	99.76	56.39	63.73	9.89
Baseline w/ Classical Acoustic Feature	72.79	69.95	61.54	75.20	60.75	75.81	61.14
Baseline w/ WavLM Embeddings	85.89	81.01	56.25	90.63	81.01	80.79	66.67
Baseline w/ ImageBind Embeddings	80.42	74.26	47.50	87.89	66.67	76.67	55.47
Baseline w/ ImageBind Multimodal Embeddings	84.88	78.48	71.25	82.16	67.05	84.86	69.09
Concat WavLM, ImageBind, and Wav2Vec2	89.49	82.28	75.00	85.99	73.17	87.10	74.07
Concat all 4 feature sets	87.91	78.82	61.54	89.60	78.69	78.87	69.06
Fusion w/ all SSL embeds projected to 3rd dim	89.84	81.01	70.00	86.62	72.73	85.00	71.33
Fusion w/ Multimodal ImageBind projected to WavLM	89.02	84.38	77.50	87.89	76.54	88.46	77.01
Fusion w/ WavLM projected to Multimodal ImageBind	89.77	85.65	73.75	91.72	81.94	87.27	77.63
Fusion w/ WavLM projected to ImageBind	88.94	85.65	75.00	91.08	81.08	87.73	77.92
Fusion w/ WavLM projected to ImageBind (Opt. Threshold)	88.94	85.65	78.75	89.17	78.75	89.17	78.75
Fusion w/ WavLM projected to ImageBind (PCA Applied)	82.22	75.95	40.00	94.27	78.05	75.51	52.89
Fusion w/ WavLM projected to ImageBind (w/ SMOTE)	74.45	73.42	58.66	86.67	64.47	76.58	61.43
Fusion w/ WavLM projected to ImageBind (10-fold cv)	90.86	85.37	81.88	89.53	80.29	88.63	81.08
PD Care Facility as Test Set	82.12	74.69	71.62	82.42	91.11	53.57	75.64
Clinical Setup as Test Set	78.44	70.20	77.32	65.83	58.10	82.57	66.35

Note that the row in bold font represents our best-performing fusion model. A more detailed performance table is reported in the supplementary material.

Table 4 | Statistical analysis across demographic subgroups

Demographic Property	Groups	# of samples	Group Accuracy	p-value	Significant?
Sex	Male	115	84.3	0.5847	No
	Female	122	86.9		
Ethnicity	White	183	82.5	0.0834	No
	Non-White	18	100.0		
Age	Below 50	19	100.0	0.1423	No
	50 and Above	199	85.9		
Recording environment	Home Recorded	133	91.0	0.0010	Yes
	Clinical Setup	72	72.2		
Recording environment	Home Recorded	133	91.0	0.9211	No
	PD Care Facility	32	93.8		
Recording environment	Clinical Setup	72	72.2	0.0176	No
	PD Care Facility	32	93.8		

In addition to the statistical bias analysis, we conducted an extended error analysis using the Microsoft Error Analysis Framework⁵⁶ to assess the reliability and fairness of our fusion model across different demographic groups. This analysis, visualized through hierarchical decision tree maps and heat maps, identified specific cohorts with elevated error rates, indicating areas where the model’s performance could be suboptimal.

From the decision tree visualization, we looked for cohort nodes with a stronger red color (representing a high error rate) branches with a higher fill line (representing high error coverage), and possible demographic combinations leading to errors. We identified a noteworthy control group of 26 individuals (Fig. 4(a)) aged above 68.5 with a significant error rate and error coverage both of 30.77%. For its all-female subgroup (14 individuals), the model exhibited the lowest performance, with an error rate of 42.86% and an error coverage of 23.08%. Besides this, we observed another PD cohort of 31 individuals (Fig. 4(b)) aged below 68.5 with an error rate of 35.48% and an error coverage of 42.31%, possessing a perfect PPV of 100% but low

sensitivity of 65%. For its all-male subgroup (20 individuals), the error rate increased to 40% with an error coverage of 30.77%. These findings suggest that model adjustments towards these specific combinations of demographics could enhance accuracy and reduce misclassifications.

Our analysis highlighted significant age-related and sex-related disparities in error rates. From the heatmap visualization (detailed in the Supplementary Note 2 and Supplementary Fig. 1, 2), we analyzed specific combinations of demographic features prone to higher error rates. Notably, white males aged 72.2–79.1 (21 participants) and 51.5–58.4 (26 participants) exhibited error rates of 33.33% and 26.92%, respectively. PPV (67% and 57%) and sensitivity (73% and 50%) scores were low for these groups. Performance was suboptimal among females as well in these age ranges (error rates 60% and 22.22% respectively), indicating critical areas for model improvement.

The model only erred for white individuals’ classifications, presumably due to the small sample sizes of other ethnicity subgroups. White males (72 individuals) exhibited a higher error rate of 18.06% compared to white

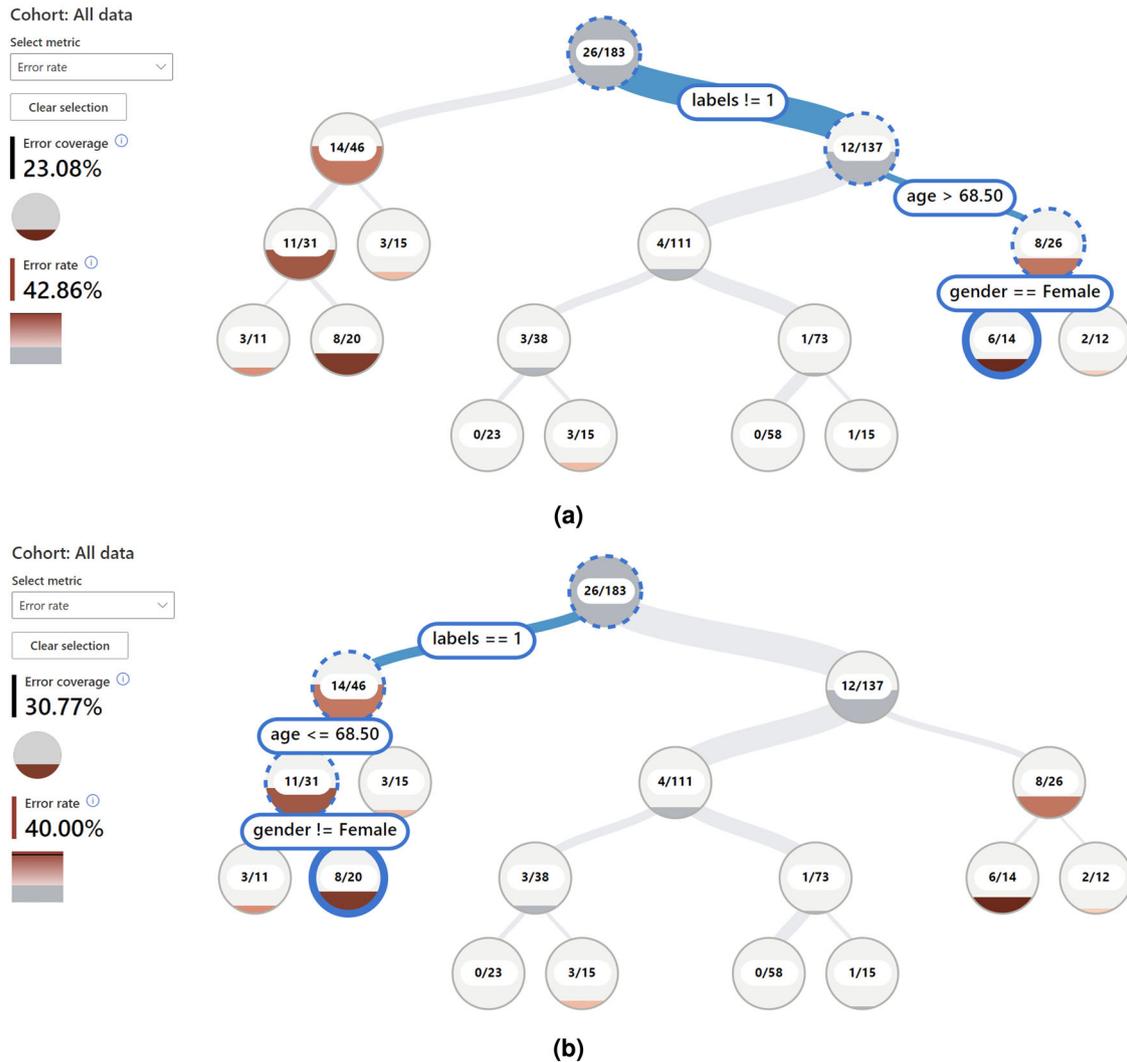


Fig. 4 | Decision tree maps of error rates and error coverage among demographic cohorts. **a** and **b** respectively demonstrate the notable nodes/cohorts with relatively high error rates and their error coverage percentage. The two numbers within each tree node represent the misclassified counts and the total counts of individuals in that

specific cohort (i.e., 26/183 indicates 26 out of 183 individuals were misclassified). Labels on the branch (i.e., age ≤68.50) represent the decision boundary condition to split the child subtrees.

females’ (93 individuals) error rate of 13.98%, suggesting possible biases in the model.

We also examined the error differences between PD and control individuals across demographic subgroups. Notably, the model performed perfectly for a large control cohort (44 individuals) at an age interval of 58.4–65.3 without any mistakes. The model erred more frequently for PD individuals aged between 51.5–58.4 (error rate = 50%, error coverage = 15.38%). Control individuals were mostly misclassified for ages between 51.5–58.4 (error rate = 15.79%, error coverage = 11.54%) and 77.2–79.1 (error rate = 36.36%, error coverage = 15.38%). PD white individuals were more often misclassified (error rate = 30.43%, error coverage = 53.85%) compared to control white (error rate = 10.08%, error coverage = 46.15%). Lastly, from the confusion matrix of sex versus PD labels, male PD individuals shared a slightly higher error rate of 32.14% and error coverage of 34.62% compared to female PD individuals’ 27.78% error rate and 19.32% error coverage. Overall, for PD classifications, the model had an error rate of 30.43% and contributed to 53.85% of the total errors, whereas for control cases the error rate is only 8.76%. This indicates a need for enhanced PPV in detecting PD, as the PD group accounted for a majority of the misclassifications. The model appeared to be biased towards avoiding false positives. Therefore, improving model sensitivity and specificity for PD

patients is crucial for reducing overall error rates and enhancing diagnostic accuracy.

In conclusion, these findings highlight the necessity for targeted interventions to address demographic-specific performance issues. This includes refining the model to reduce biases, enhancing training data diversity, and implementing demographic-specific adjustments to improve accuracy and fairness across all groups. The implications of these error patterns are critical for developing a more robust and equitable AI-driven healthcare solution.

Ablation Studies

Our ablation studies thoroughly evaluated the effectiveness of various feature sets for PD classification. We adopted different combinations of features sets for concatenation and also fused after projection to boost the metrics. For a detailed presentation of these results, please refer to the Supplementary Note 1 and Supplementary Table 1.

Discussion

In today’s digital age, mobile devices have become pervasive across global populations, encompassing all age groups⁵⁷. These devices universally feature capabilities for audio recording, providing a practical platform for

deploying our speech-based PD screening framework. By simply reciting a standard pangram, users can leverage their mobile devices to conduct preliminary screenings for PD. Our research further offers the potential to develop a mobile application utilizing semi-supervised speech models and fusion architecture and could continuously analyze natural speech during phone conversations—with explicit user consent—to detect early signs of PD and generate timely alerts. Such technological advancements significantly diminish the necessity for frequent clinical visits, offering a substantial benefit to individuals in areas where access to specialized neurological care is limited. This approach not only facilitates convenient at-home monitoring but also plays a crucial role in the early detection and managing the progression of PD, potentially altering the course of the disease by enabling earlier therapeutic intervention.

In 2017, PD placed a significant economic burden of \$52 billion in the United States⁵⁸. Given the projected doubling of PD patients by 2030⁵, the scenario will worsen significantly, surpassing \$79 billion even without accounting for inflation⁵⁸. This growing financial strain underscores the crucial need for early diagnosis and regular monitoring, which have been shown to significantly improve patient outcomes and reduce the overall burden on healthcare⁵⁹. Minimizing unnecessary clinical visits not only offers potential for substantial cost savings for both patients and healthcare providers but also means that each avoided visit—which might otherwise be spent assessing healthy symptoms—can save patients significant amounts on consultation fees and associated travel expenses. The introduction of remote screening models like ours could therefore transform the management of PD, leading to a more efficient allocation of healthcare resources and financial savings.

Dashtipour et al.⁶⁰ highlighted that speech impairment affects up to 89% of individuals with PD. In contrast to methods that require sustained phonation, the analysis of free-flow speech provides a more natural and comprehensive assessment of vocal impairments. This approach captures a wider spectrum of vocal characteristics and abnormalities, offering the potential for more accurate and earlier diagnosis than is possible with phonation-based models alone. Furthermore, the assessment of natural speech forms a core component of in-person evaluations as outlined by the MDS-Unified Parkinson's Disease Rating Scale (MDS-UPDRS)¹, whereas sustained phonations are not explicitly monitored under these guidelines. While our study did not directly analyze continuous speech, the use of pangram utterance closely approximates natural speech patterns. The semi-supervised models employed in this research, which are trained on natural speech, are thus well-suited to capture the speech dynamics indicative of PD, even from the structured utterance of a pangram.

The projection-based fusion model has demonstrated superior performance compared to simple concatenation. The model projecting WavLM features into the ImageBind dimension achieved an AUROC of 88.94% and accuracy of 85.65%, outperforming models using concatenated features, which despite a slightly higher AUROC of 89.49%, had lower accuracy at 82.28%. This result implies that projection-based fusion effectively aligns and synergizes different feature sets, overcoming issues like redundancy and scale disparity commonly seen with concatenation. Notably, this fusion method also surpasses models built on hand-crafted acoustic features, highlighting the limitations of traditional feature engineering in capturing the subtle, discriminative nuances of PD in speech. In contrast, embeddings derived from semi-supervised models like WavLM and ImageBind probably revealed complex PD-related speech patterns, suggesting that automated, embedding-based features significantly enhance the precision and depth of PD speech analysis – a key contribution of our study.

Rizzo et al.⁶¹ reported that the accuracy of PD screening by non-expert clinicians stands at 73.8%, with a 95% credible interval (CrI) of 67.8%–79.6%. In contrast, movement disorder specialists achieve a slightly higher accuracy of 79.6%, albeit with a wider CrI of 46%–95.1%. Encouragingly, our model demonstrates promising results, achieving an accuracy of 85.65%, positioning it favorably within or even above these established confidence intervals. However, it is important to note that the

populations used to evaluate our model and those in clinical studies are very different, so while this comparison is promising, it should be interpreted with caution. Our model's performance, coupled with its generalization capability across diverse recording environments, underscores its potential for global deployment among English-speaking populations. When tested on datasets from a PD care facility and a clinical setup, which were completely unseen during the training phase, our model demonstrated respectable AUROCs of 82.12% (accuracy of 74.69%) and 78.44% (accuracy of 70.20%), respectively. These results are comparable to, and in some cases surpass, those achieved by non-expert clinicians, highlighting the model's ability to deliver reliable PD screening. The observed drop in metrics during external validation can be attributed to variations in recording devices and environments, changes^{62,63}. However, this challenge also highlights a key strength of our approach compared to existing models, which are typically trained on data from a single recording device. A distinctive characteristic of our home-recorded dataset is the diversity in recording devices, which mirrors real-world conditions better than controlled datasets. Nevertheless, we acknowledge that substantial differences in recording equipment can potentially degrade the model's performance. To mitigate the potential risk of overfitting to specific training environments and to ensure patients' safety, we recommend collecting a few initial data samples from the local environment at new deployment sites and retraining the model for optimal performance. In the future, we aim to further enhance the model's generalizability across diverse and unpredictable settings by continuously incorporating more data from unseen environments and retraining the model on increasingly heterogeneous data sources. While this study does not explore continual learning⁶⁴, we plan to apply various continual learning algorithms in the future to effectively address the challenge of integrating new incoming data, ensuring the model remains adaptable and robust over time.

Our statistical significance tests show that the model demonstrates broad invariance to demographic diversity, with no significant bias detected across sex, ethnicity, or age subgroups. While we observed significant changes in performance across different recording environments, this finding further supports our objective of enhancing the model's generalizability by incorporating data from varied recording setups. We also recognize that the analysis between the white and Non-White subgroups was limited in scope due to insufficient sample sizes of each of the Non-White ethnic groups. As a result, we were compelled to group all Non-White participants into a single category, which prevents a nuanced understanding of the model's performance across granular ethnic subgroups. This limitation affects the generalizability and reliability of the model for Non-White populations, leaving the findings somewhat inconclusive. Notably, this disparity in demographic representation, particularly within PD datasets, is a known challenge in the literature. For instance, nearly 90% of the participants in the PPMI dataset⁶⁵ belong to the White ethnic group. Moving forward, we aim to address this limitation by collecting more data from underrepresented demographic subgroups and retraining the model incrementally to ensure optimal performance across diverse populations.

Despite the model's overall robustness across demographic groups in terms of hypothesis testing, our detailed error analysis identified certain under-performing data cohorts with notable misclassification rates, indicating that we should be cautious interpreting the model's prediction results upon these cohorts. While the model performed well for males under 51 and over 72—likely due to more distinct PD traits in these age brackets—the transition age of 51 to 72 presents challenges. The speech patterns in this middle age group are less distinct, possibly due to physiological changes that are harder for the model to interpret. Subtle changes in speech in this age range may reflect both PD-related traits and other aging factors, making it harder for the model to distinguish the true source of variation. For females, the group over 72 also presents higher error rates, with 4 out of 6 errors being cases where the model predicted Non-PD participants to have PD. This misclassification may be due to significant physiological changes in vocal mechanisms, such as the thickening of the epithelium post-70⁶⁶, which alters vocal characteristics and potentially confuses the model. The changes in

speech patterns at higher ages for females might mimic anomalies associated with PD, causing the model to incorrectly identify them as PD cases. Studies support these observations, noting that males experience marked structural changes in their vocal mechanisms around age 60, while females undergo notable changes post-70. Males exhibit an increase in spectral energy skewness and nonlinear changes in fundamental frequency (F0) with age, whereas females show nonlinear changes in signal-to-noise ratio (SNR), further complicating model predictions in these age and gender groups^{66–68}. These variations in vocal mechanisms, along with sex-specific acoustic properties, likely contribute to the increased error rates observed in certain demographic groups in our study. Supporting this, our SHAP-based feature importance analysis reveals that the embedding dimensions most influential in the misclassifications of these subgroups are significantly correlated with age, reinforcing the hypothesis that age-driven physiological changes impact model predictions and contribute to the observed misclassification patterns (see Supplementary Note 2 for detailed feature importance and correlation analysis). Future studies should further explore these biological and neural factors in speech-based PD predictions to achieve a more in-depth understanding. Moreover, variations in misclassification linked to ethnic differences in speech patterns could be explored with a more ethnically diverse dataset, enhancing our understanding of the model's demographic discrepancies. The overall sensitivity of our model is relatively lower at 75.0%, and it drops further to 70.0% among the 46 PD participants with available demographic data. One of the contributing reasons for this higher false negative rate could be the imbalance in our dataset, where the number of PD participants is significantly lower than the Non-PD group (another potential factor might be the model's sole reliance on speech, as discussed later in this section). While data resampling techniques were considered to mitigate the data imbalance issue, the hyper-parameter tuning approach revealed that even effective approach, such as SMOTE was insufficient to improve the predictive performance, which coincides with previous findings that the higher dimensionality^{51,69} of embeddings from models like WavLM and ImageBind can limit the efficacy of such techniques. Although recruiting PD participants poses challenges due to logistical constraints, lower patient availability, and the complexities of remote data collection, in future, we plan to collaborate closely with clinics and research institutions, facilitating access to larger PD cohorts and ensuring a more balanced and representative dataset. In the meantime, we advise cautious interpretation of the model's results to minimize any potential risks from misclassification, particularly with the higher false negative rates.

One of the principal limitations we must acknowledge is the limited explainability of our model. While the use of vector embeddings from semi-supervised models like WavLM provides a powerful tool for feature extraction, the black-box nature of these models poses challenges in interpretability. While tools like SHAP⁷⁰ or LIME⁷¹ were considered for improving explainability, the abstract nature of embeddings—where dimensions lack inherent meaning—limits their utility. At best, these tools could highlight which part of the embedding space contributes most to predictions, but this insight remains too abstract for clinical relevance. We recognize the need for greater transparency of our model and plan to explore more suitable methods for model explainability in future work (potentially through novel approaches), as the field of explainable AI for embeddings is still evolving^{72,73}. Additionally, our current reliance on English pangrams restricts the model's applicability to non-English speakers. However, the inherent adaptability of semi-supervised speech models offers a promising avenue for extending our approach to other languages. Recent research has shown significant progress in applying semi-supervised transfer learning to expand speech recognition models for low-resource languages. Recent works have demonstrated the application of semi-supervised models for low-resource languages, such as methods for Italian⁷⁴ and aphasic speech recognition in English and Spanish⁷⁵. Advances like using self-supervised representations for multilingual language diarization⁷⁶ further illustrate the growing potential for handling multilingual and low-resource languages. In future work, our group aims to explore fine-tuning or adapting existing models to generate speech embeddings for diverse languages. Given our

model's relatively simple architecture and low-resource training requirements, once these models for other languages are available, we will be able to adapt and extend our approach to non-English speakers. This future direction not only highlights the potential for our model to accommodate speakers of different languages but also emphasizes the broader evolution of semi-supervised speech models beyond linguistic boundaries. As an preliminary analysis of our model's extendability, we took our best model trained only with pangram utterance data and tested it on a completely separate test dataset where the participants delivered a continuous, free-flowing speech for around one minute on their preferred topic (not any pangram). This test dataset involves 177 participants, and 39 of them are diagnosed with PD. Our model, even without seeing any such continuous speech, achieved a respectable AUROC of 77.4% with an accuracy of 74.1%. This underscores the potential versatility of our approach in broader speech analysis contexts, not just for PD detection but possibly for other speech-related applications as well. In addition to the model's reliance on English language, a further challenge affecting its global applicability is the limited access to reliable technology, such as desktop or laptop computers with stable internet connections. Although addressing this constraint is beyond the scope of this study, we suggest that in regions where such resources are scarce at an individual level, deploying the tool at community-accessible locations equipped with necessary technology could help minimizing this barrier to some extent. While this setup may not fully achieve the level of accessibility we envision, it could nonetheless improve the tool's usability in underserved, remote areas.

The symptoms of PD vary widely among individuals, with some primarily exhibiting significant speech impairments, while others may experience motor symptoms like tremors without noticeable changes in their speech patterns. Consequently, our model, which primarily analyzes speech dynamics, may not be universally effective for all PD patients, especially those whose vocal symptoms are less pronounced early in the disease. This could also contribute to the model's higher false negative rate. To address this, we recognize the potential of integrating additional modalities into our framework. Given the flexibility of being a web-based platform, incorporating video data collection for tasks like finger tapping (for motor assessment of bradykinesia⁷⁷) and facial expressions (for hypomimia⁷⁸) is feasible. Moreover, our preliminary experiments with multimodal embeddings generated by ImageBind suggest promising potential for leveraging both audio and video data in a unified analysis. Moving forward, we plan to incorporate these modalities into a unified model that considers speech, motor function, and facial expressions, offering a more comprehensive and reliable PD screening tool. This multimodal approach will allow the model to capture a broader spectrum of PD symptoms, improving the overall performance and utility across diverse patient groups. In addition, validating the model's robustness across different stages of PD and analyzing the impact of MDS-UPDRS assessment scores for the speech task on model performance would have provided further insights. However, collecting PD stage data and UPDRS scores in a home-based data collection procedure presents significant logistical challenges. Gathering this data requires substantial time and effort commitments from specialists, which we acknowledge as a limitation of the study. Moving forward, we plan to conduct a controlled data collection process where clinicians can be involved to provide UPDRS assessments with staging of the disease, allowing us to explore the relationship between speech-based model performance and disease severity more thoroughly.

Integrating AI into healthcare brings significant ethical challenges, particularly regarding data privacy and the accuracy of assessments. To ensure data privacy, any deployed version of the tool should remove video and audio data immediately after feature extraction to mitigate unauthorized access risks, especially in large-scale clinical or at-home deployments. In a real-world scenario, another concern is the potential for inaccurate risk assessments, as false positives, observed in 8.92% of cases, may cause unnecessary anxiety but could also prompt users to seek professional medical advice. To alleviate the distress, we recommend that the deployed system clearly inform users that the results are not a definitive diagnosis and

not free from errors. Furthermore, integrating psychological care resources, such as support groups, would ensure users have access to emotional support if distressed by their results. Conversely, the 25% false negative rate is more concerning as it could delay the necessary medical care. As we already discussed, this limitation may be due to dataset imbalance, as data from PD patient is harder to collect or the model's sole reliance on speech task, or perhaps both. Moving forward, we aim to address this by balancing the dataset and incorporating additional tasks, potentially enhancing model performance. Meanwhile, users must be informed of the tool's limitations on misclassification, emphasizing the importance of consulting healthcare providers regardless of the device's output.

The implications of this study are broad, extending beyond PD diagnosis. The methodologies developed could be adapted for identifying other speech-related deficiencies, offering a blueprint for future research in neurological disorders. By integrating our proposed projection-based fusion architecture with semi-supervised speech embeddings, we anticipate similar methodologies could significantly improve the performance of models in various speech analysis applications. The success of this project highlights the transformative potential of AI in healthcare, particularly in enhancing diagnostic processes through advanced machine learning techniques and accessible digital platforms.

Methods

Dataset Description

For collecting the speech dataset used in this study, we employed the web-based PARK framework developed by Langevin et al.⁴⁶, accessible at <https://parktest.net/>. Participants worldwide can use this framework to record themselves while performing tasks inspired by the MDS-UPDRS guidelines designed to assess motor symptoms for evaluating PD. One such task involved the articulation of a standard English pangram, “quick brown fox.” To ensure consistency and proper execution across participants, the full pangram text was visible on the web interface for them to read aloud, and an instructional video was provided before each task. Additionally, participants were required to complete questionnaires that captured demographic information such as age, sex, ethnicity, PD diagnosis year (for participants with PD), etc. From the video recordings, we extracted audio clips to compile our dataset. Due to some participants providing multiple video/audio samples at different times, we could amass a total of 1854 samples for our study. Additionally, to assess the extendability of our model to more natural speech scenarios from a pangram-only setting, we conducted a supplementary test. In this test, 177 participants, including 39 with PD, were instructed to speak freely for one minute on a topic of their choice. This additional dataset allows us to explore our model's performance in continuous free-flow speech settings, further underlining its potential applicability to broader speech analysis contexts. Note that this study was approved by the Institutional Review Board (IRB) of the University of Rochester and the University of Rochester Medical Center. Informed consent was collected electronically due to the remote nature of the study, authorizing the use of participants' data and images.

We collected the dataset from three distinct settings.

- **Home Recorded:** We gathered a significant portion of our dataset from participants who recorded themselves staying at home using the PARK tool. We reached these participants by advertising our PARK tool on social media. We also emailed individuals who were willing to contribute to PD research. Despite being a global effort, we could only collect data from 67 (10% of this cohort) PD participants. In this data collection setup, the labels of the participants (PD or control) were self-reported.
- **Clinical Setup:** In collaboration with the University of Rochester Medical Center (URMC) in New York, participants in a clinical study recorded themselves using the PARK tool. This setting ensured some supervision by clinical staff, particularly for those participants who required assistance during the recording. The participants of this cohort were clinically confirmed to be PD or control. Almost 30% of the total PD participants are from this cohort.

- **PD Care Facility:** Our last data collection site was InMotion (details at <https://beinmotion.org/>), a PD care facility in Ohio, US. We could collect the video/audio data from their clients as clinically confirmed PD patients and their caregivers as self-reported controls. This environment provided a supportive setting for participants, typically involving assistance from their caregivers and/or the InMotion's staff during recording sessions. We collect the major portion of PD samples (47%) from this setup. Note that the supplementary dataset of continuous free-flow speech involving 177 participants, including 39 with PD, was also gathered here.

In this study, we collected data across a broad spectrum of demographic groups, targeting a comprehensive analysis of PD across diverse populations. The dataset featured a balanced sex distribution with 53.2% female and 46.6% male participants but showed a predominance of white participants at 66%, with 25% not disclosing their ethnicity. Participants ranged from 16 to 93 years old, with a majority aged 60 – 69 years, which is significant as PD prevalence increases with age. Data collection occurred in varied settings to enhance external validity: 49.9% at home, 27% in clinical setups, and 20.7% at a PD care facility, reflecting the accessible nature of our methodology. We collected PD stage information from 213 study participants—105 out of 117 from the Clinical Setup and 108 out of 185 from the PD Care Facility—based on their clinical records. The distribution of PD stages was relatively balanced, ranging from stage 0 (Asymptomatic) to stage 3 (Moderate symptoms with postural instability, requiring assistance to recover from the pull test). In addition to PD stages, we obtained disease duration data from 143 participants, with more than half having lived with PD for less than five years, a period generally considered the early phase of disease progression.

To enhance the quality and utility of our dataset, we implemented a comprehensive data cleaning and preparation process for videos captured under varied conditions. Initially, we standardized the video format by converting all WEBM files to MP4, ensuring consistent fps and uniform metadata. We then isolated the audio from each video, converting it to a WAV format sampled at 16 kHz, which is optimal for detailed acoustic analysis. Using the Whisper To precisely capture segments containing the target pangram, we used the Whisper model⁷⁹, which provided robust transcription and timestamping. This process involved detecting occurrences of specific keywords within the pangram – ‘quick’, ‘brown’, ‘fox’, ‘dog’, and ‘forest.’ Whisper generated start and end timestamps for each identified segment, and we defined each clip using the start timestamp of the first keyword and the end timestamp of the last keyword in the detected sequence. We also extended each audio segment by 0.5 seconds on either end to retain contextual audio cues, such as breaths, subtle background noise, and speech transitions, which provide valuable context around speech boundaries and contribute to a more informative representation of natural speech patterns. These extended segments were then saved in WAV format to create a comprehensive speech dataset for further analysis.

Digital Speech Feature Extraction

This study aims to objectively and quantitatively capture the nuanced speech dynamics that can be pivotal in differentiating PD characteristics. To achieve this, we extracted a series of classical acoustic features alongside advanced deep learning embeddings using state-of-the-art models such as Wav2Vec2³⁷, WavLM³⁶, and ImageBind³⁸. The remaining part of this subsection details the methodologies employed for each type of feature extraction.

We extracted classical acoustic features proven to be crucial in the literature in characterizing speech disorders associated with PD^{6,15,19}. These features include Mel-frequency cepstral coefficients (MFCCs)^{16,17}, jitter^{16,18}, shimmer^{18,19}, and pitch-related metrics^{20,21}:

- MFCCs represent the short-term power spectrum of sound, based on a linear cosine transform of a log power spectrum on a nonlinear mel scale of frequency⁸⁰.

- Jitter measures frequency variations from cycle to cycle, offering insights into the stability of vocal fold vibrations⁸¹.
- Shimmer quantifies amplitude variations, useful in assessing vocal fold closure inconsistencies⁸¹.
- Pitch feature encapsulates the fundamental frequency, providing crucial information on the tonal aspects of the speech⁸².

Tools such as Praat^{83,84} and Parselmouth⁸⁵ were utilized to calculate these features from the digitized voice recordings of the participants.

One of the significant contributions of our study is the use of deep embeddings of the speech audio extracted by three distinct pre-trained semi-supervised language (SSL) models: Wav2Vec 2.0 (W2V2), WavLM, and ImageBind. Our selection of these models was motivated by their robustness in capturing nuanced audio features and their ability to handle noisy speech data in complex audio environments, which is especially relevant for detecting subtle speech variations associated with PD.

Wav2Vec 2.0 (W2V2)³⁷ was developed by Facebook AI, which utilizes a self-supervised learning framework to learn high-quality representations from raw audio waveforms. By predicting masked segments within the audio context, Wav2Vec 2.0 captures essential speech dynamics and produces embeddings well-suited for identifying subtle acoustic patterns, which is particularly valuable for PD detection.

WavLM³⁶ from Microsoft was built upon Wav2Vec 2.0's architecture, which enhances the model's ability to handle diverse acoustic environments, including noisy backgrounds and overlapping speech, making it highly effective for applications in real-world scenarios. This model's ability to adapt to unpredictable audio environments makes it effective for clinical or home-based PD detection, where environmental consistency cannot always be controlled, and its capacity to discern subtle audio patterns in noisy settings aligns well with the clinical requirements of PD speech analysis, offering refined acoustic detail extraction critical for our task.

ImageBind³⁸, introduced by Meta AI, extends beyond conventional audio-only models by incorporating cross-modal learning between audio and visual data. By training on paired datasets of images and corresponding audio, the model learns to create embeddings that reflect not just the audio content but also its relation to visual elements, enhancing the ability to discern nuanced variations in speech possibly linked to neurological conditions like PD.

Our decision to focus on these specific models over alternatives like UniSpeech⁸⁶, HuBERT³⁹, and XLS-R⁸⁷ stems from their particular optimization goals. While these models are recognized for their strengths in multilingual and cross-lingual tasks, their emphasis on Automatic Speech Recognition (ASR) and speaker recognition does not fully align with our objective of extracting helpful acoustic high-resolution embeddings for PD detection where robustness to noise and subtlety in speech variation is more critical.

To obtain deep embeddings, we fed each extracted pangram audio sample into our selected SSL models. These models process the raw audio input through multiple layers, progressively encoding complex representations of the speech data. We extracted the embeddings from the final hidden layer, capturing the most sophisticated, context-rich features of the audio. This focus on high-resolution, contextually robust embeddings provides a solid foundation for our analysis, empowering our approach to capture the nuanced acoustic anomalies linked with PD.

Feature Pre-processing

Before starting the training phase, several data pre-processing steps were employed to optimize the dataset for effective machine learning model training.

Data Deduplication was performed to ensure dataset integrity by removing duplicates based on data collected from the same participant on the same date, ensuring that each speech sample was unique and properly represented in the dataset. This process eliminates redundant observations, prevent any potential bias, maintains data quality, and avoids over-representation in model training.

Correlation Analysis uses Pearson correlation coefficient⁸⁸ to generate a correlation matrix among the feature set, and features exhibiting a correlation coefficient above a predefined threshold were identified and elimination. This process mitigates multicollinearity, potentially enhancing model stability and interpretability while reducing dimensionality. The threshold and the decision to drop correlated features were set as tunable parameters, allowing for flexibility and optimization based on the specific characteristics of our dataset. Note that after hyperparameter tuning, the best model eliminated correlated features when the coefficient exceeded a threshold of 0.85.

Data Scaling is essential to prevent features with larger scales from dominating model training, ensuring each feature contributes proportionally⁸⁹. In this study, we explored two popular scaling methods—Min-Max Scaling and Standard Scaling—as they are well-suited for neural network-based architectures. Min-Max Scaling, which maps features to a specified range (typically [0, 1]), is effective for features with varying ranges, promoting faster convergence by keeping all values within a controlled scale⁹⁰. On the other hand, Standard Scaling, on the other hand, centers features around a mean of 0 with a standard deviation of 1, making it ideal for algorithms that benefit from normally distributed data or are sensitive to feature variance⁹¹. We configured the choice of scaling method (including using it or not) as a hyperparameter, allowing optimization based on data characteristics and model requirements. Min-Max Scaling is particularly useful when there are large disparities in feature ranges, while Standard Scaling is preferred when the data distribution approximates Gaussian^{92,93}. Applying these scaling techniques contributes to stable training, especially for gradient-based algorithms, by reducing distributional disparities across training, validation, and test sets, thereby enhancing model robustness and generalizability⁹⁴. After hyperparameter tuning, Standard Scaling was selected for the best model performance.

Data Resampling was employed to address the dataset imbalance, as the proportion of PD samples was approximately half that of control samples across all cohorts. Imbalanced datasets can lead to biased models that overfit the majority class and underperform in predicting the minority class, which, in this case, was the PD class. To mitigate this, we applied three resampling techniques: the Synthetic Minority Over-sampling Technique (SMOTE)⁵¹, which generates synthetic samples from the minority class; Random Undersampling⁵², which reduces the control sample count to match the PD class; and Random Oversampling⁵², which duplicates minority class samples to achieve balance. We configured the choice of using one of these three resampling techniques alongside the option of not using any resampling as a tunable parameter, allowing a systematic evaluation of their impact on model performance. However, after hyperparameter tuning, none of these techniques contributed to improved performance in the best model. As such, exploring and implementing more advanced data balancing techniques remains an area for future work, aiming to further mitigate the potential effects of data imbalance on model accuracy and robustness.

Baseline Modeling

Following the setup of our data pre-processing and model development pipeline, we conducted initial baseline experiments. Our pre-processing pipeline, designed to be feature agnostic, enabled independent training of models on different feature sets (Classical, W2V2, WavLM, ImageBind). To benchmark performance, we tested three primary model architectures: a Convolutional Neural Network (CNN) trained directly on raw speech data, Support Vector Machine (SVM) classifiers trained on extracted classical and deep embeddings, and deep learning models with fully connected layers also trained on these extracted feature sets. For the CNN model, we designed an end-to-end PD classification pipeline aiming to capture PD-specific characteristics directly from time-series data. In addition, we used SVM classifiers to explore linear and non-linear decision boundaries within each feature set.

For our neural-network based baseline models, we employed two structures: a shallow fully connected classification layer (ShallowANN) with sigmoid activation and an Artificial Neural Network (ANN) with an

additional hidden layer before the output layer. During the training of DL models, the choice between ShallowANN and ANN was kept as a hyperparameter to optimize performance based on the dataset characteristics and feature set.

Fusion Modeling

In our study, we explored several strategies to fuse multiple feature sets, enhancing the robustness and accuracy of the resulting models. We began with a simple vector concatenation approach, where distinct feature sets were merged into a single dataset. This basic concatenation strategy allowed us to establish a baseline for further fusion methods. Using the concatenated datasets, we trained both shallow and deep neural networks (ShallowANN and ANN) to assess the performance of combined features.

Moving beyond simple concatenation, we implemented a hybrid fusion approach, termed the “*Projection-based Fusion Architecture*” in Fig. 1, which leverages projection and reconstruction techniques. Below, we describe the key components of this architecture.

Projection Layer in our model architecture is essential for aligning multi-modal features within a shared latent space, enabling effective fusion of diverse feature types such as WavLM, Wav2Vec2, and ImageBind embeddings. This approach addresses the limitations inherent in simple concatenation, which often leads to feature mismatches, added noise, and increased dimensionality that can compromise model performance, particularly in high-dimensional, nuanced data like PD-related speech patterns^{95,96}. By aligning features from different modalities, the projection layer enhances cross-modal interactions, making it possible for the model to leverage nuanced, multi-modal patterns crucial for detecting the subtle speech variations indicative of PD^{13,42}.

Our choice to incorporate a projection layer is driven by several key considerations. First, by transforming features into a shared, lower-dimensional space, the projection layer reduces dimensionality and mitigates the risk of overfitting, which is especially important given the high-dimensional nature of speech features⁹⁷. Additionally, it enhances cross-modal alignment, allowing for deeper feature interactions that are essential for capturing PD-specific speech characteristics⁴³. The layer also normalizes features across modalities, addressing scale and distribution differences that could hinder simple concatenation methods⁹⁸. This adaptive, learnable space enables dynamic feature fusion during training, promoting robust and generalizable representations – a critical asset in PD detection where symptoms are subtle and varied⁹⁹. Our approach aligns with recent advancements in multimodal learning^{100,101}, where projection layers are employed to unify diverse data types within a common space, as seen in models like the Multimodal Video Transformer (MMV)¹⁰² and Multimodal Masked Autoencoder¹⁰³. These examples underscore how projection techniques enhance joint representation learning, benefiting tasks such as action recognition and pretraining.

To optimize the projection layer, we conducted extensive hyperparameter tuning to find an ideal balance between information retention and model efficiency¹⁰⁴. We experimented with a range of projection dimensions, from compact to more expansive, and further explored projecting one feature set into the space of another. Ultimately, the model achieved optimal performance by projecting WavLM features into the ImageBind feature space, which facilitated a richer, cross-modal alignment that was particularly proven to be effective in capturing the nuanced, multi-modal speech patterns relevant to PD detection¹⁰⁵.

Fusion Layer integrates the projected and aligned features into a cohesive, unified representation. First, by normalizing the projected features, this layer reduces distributional disparities across modalities. The normalized features are then combined, either by direct summation or by summing and further normalizing them, creating a unified cross-modal representation. This approach not only harmonizes diverse modality features but also enhances the model’s ability to capture rich, multi-modal interactions, a critical aspect in tasks where nuanced multi-modal patterns are essential.

Decision Layer processes the unified representation to yield a final classification output. Two variations of the decision layer are employed as hyperparameter options: a deeper fully connected network and a simpler, shallow architecture. In the shallow configuration, the decision layer consists of a single fully connected layer, followed by a sigmoid activation, directly mapping the fused representation to a probability score. In contrast, the deeper variant introduces an intermediate layer, adding non-linearity and enabling the model to capture more intricate patterns in the fused representation before the final classification. The choice between these versions was tuned as a hyperparameter. The streamlined structure of both decision layers reduces model complexity and minimizes the risk of overfitting, particularly when working with limited data or high-dimensional features.

Loss Function for the projection-based fusion model is a multi-objective function with three components, each serving a crucial role in optimizing the model’s performance and robustness:

- Prediction Loss: This component uses Binary Cross-Entropy (BCE) to reduce the disparity between model predictions and ground truth, driving accuracy in classifying PD and non-PD cases¹⁰⁶.
- Cosine Loss: Calculated as the complement of cosine similarity between projected multi-modal features, this loss component guides the model to learn diverse, complementary representations¹⁰⁷. It aims to capture unique aspects of each modality, potentially improving the model’s ability to leverage multi-modal information.
- Reconstruction Loss: This loss ensures the fidelity of feature reconstruction from the projected space, using one chosen metric (hyperparameter) from Mean Squared Error (MSE), L1 norm, or Kullback-Leibler (KL) divergence¹⁰⁸. The primary objective is to preserve essential input information while balancing dimensionality reduction and retention.

The combined weighted sum of these loss components enables fine-tuning of each objective’s influence¹⁰⁹, optimizing classification, feature representation, and information retention for a more robust and generalizable PD detection model.

Our model architecture is intentionally designed to be relatively simple yet effective for PD detection, balancing performance with computational efficiency. Leveraging pre-trained models like WavLM, Wav2Vec2, and ImageBind for feature extraction allows our model to capture complex, high-dimensional representations without requiring additional deep layers, benefiting from transfer learning to enhance performance with limited PD-specific data^{37,110,111}. This approach reduces the risk of overfitting, which is common in complex architectures, especially with limited datasets in medical applications¹¹². By maintaining an optimal data-to-parameter ratio, we improve generalizability and ensure stable performance across data splits^{113,114}. Additionally, the simplicity of our model enables faster training and lower computational demands, aligning well with real-world clinical applications where efficiency and robustness are essential^{115,116}. Despite its streamlined design, our ANN strikes a balance between model complexity and generalizability, ensuring strong performance without unnecessary computational overhead. This is consistent with findings in prior studies, where even simple architectures—such as linear layers—have been shown to perform effectively for complex tasks like ECG-based AFib detection¹¹⁷ and vision-language modeling¹¹⁸, when paired with efficient pre-trained models to represent intricate data patterns. Inspired by efficient architectures like LLaVA, which successfully integrate multi-modal data, our model proves that high accuracy and stability can be achieved without excessive complexity¹¹⁸.

Evaluation Metrics

To assess the performance of our deep learning models effectively, we used several key metrics that are crucial for clinical evaluation: Area Under the Receiver Operating Characteristic (AUROC) score, Accuracy, Sensitivity, Specificity, Positive Predictive Value (PPV), and Negative Predictive Value (NPV). We visualized the model’s performance through AUROC curves

and confusion matrices. The AUROC curve visually illustrates the model's discrimination ability between different conditions, and the confusion matrix details the counts of true positives, false positives, true negatives, and false negatives, essential for evaluating the model's practical efficacy. The model achieving the highest AUROC score on the validation set was selected and then tested on the test set with these metrics to confirm its clinical relevance and effectiveness.

Statistical Bias Analysis

To ensure equitable performance across diverse populations, we conducted a comprehensive bias analysis across demographic subgroups, specifically based on age, sex, ethnicity, and recording environment. Participants without demographic data were excluded from this analysis to maintain data integrity. Given the multiple comparisons involved, we performed six distinct hypothesis tests and applied a Bonferroni correction⁵³ to control for Type I error, adjusting our effective significance level to $\alpha^* = \frac{0.05}{6} = 0.0083$.

For each comparison, we used Fisher's Exact Test to assess statistical significance. This test was chosen for its robustness with categorical data and its suitability for our sample size, as it does not require the normality assumption of parametric tests like the Z-test⁵⁵. Fisher's Exact Test provides an exact *p*-value, making it appropriate for detecting differences in performance proportions between subgroups without the reliance on large sample conditions.

The bias analysis included subgroup comparisons by sex (Male vs. Female), ethnicity (White vs. Non-White), age (Below 50 vs. 50 and Above), and recording environment (Home Recorded, Clinical Setup, PD Care Facility). For recording environment, we conducted pairwise statistical tests between each setting to provide a detailed evaluation of how recording conditions may impact model performance. To further examine potential biases, we also analyzed the influence of disease duration on model performance by calculating Spearman's rank correlation among participants with known disease durations. This non-parametric test was selected to measure the association between model accuracy and disease duration without assuming a linear relationship. To examine the association between PD stages and model accuracy, we employed both the Kruskal-Wallis H-test and Fisher's Exact Test with a Monte Carlo simulation approach (10,000 replications) to estimate *p*-values. The Kruskal-Wallis H-test is a non-parametric statistical test used to determine whether there are significant differences between multiple independent groups. It is particularly suitable when the assumption of normality is not met, as it compares the median ranks across different PD stage categories. This test helps assess whether model accuracy varies significantly across PD stages. Additionally, we applied Fisher's Exact Test with Monte Carlo simulation, as it is well-suited for small sample sizes within certain PD stage groups. Since our dataset included four distinct PD stages, Fisher's Exact Test does not yield a single odds ratio. However, the Monte Carlo simulation method enhances reliability by empirically generating the distribution of the test statistic, allowing for robust *p*-value estimation despite the sparsity of data in certain subgroups.

Cohort-based Error Analysis

Building on the statistical bias analysis which assessed model fairness across demographics, our cohort-based error analysis delved deeper into individual and group-level performance nuances. This fusion model utilizes deep learning architecture and fusion features, which generally offer limited intuitive interpretability and explanation of the model's decisions, as well as challenges in identifying when and why the model errs. Post-assessment using AUROC and other standard metrics, we tracked both predicted and true PD labels to explore the model's operational dynamics under varying demographic influences. We employed Microsoft's Error Analysis SDK⁵⁶ to visually dissect error patterns and concentrations, particularly focusing on cohorts with heightened error occurrences. The error analysis process involved:

- Cleaning the data by removing entries with missing demographic information (age, sex, ethnicity) and consolidating the test set to

include necessary true and predicted labels and demographic features for each participant.

- Employing a decision-tree-like hierarchical structure to systematically identify error instances and understand prevalent error patterns by demographic features.
- Segregating the data into cohorts based on combinations of demographic characteristics to explore correlations with high prediction errors. Error rates, coverage, PPV, and sensitivity were visualized through heat maps for combinations – {sex × ethnicity}, {sex × age}, {age × ethnicity}, {sex × true PD labels}, {age × true PD labels}, {ethnicity × true PD labels}.

This structured approach involving decision tree maps and heat maps allowed us to pinpoint under-performing demographic subgroups and potential triggers, enhancing our understanding of the model's performance across diverse populations.

Hyperparameter Tuning

Hyperparameter tuning^{119,120} is crucial for optimizing machine learning models to enhance their predictive capabilities. In our study, we rigorously tuned both the Baseline Modeling and Fusion approaches with the goal of maximizing the Area Under the Receiver Operating Characteristic (AUROC) for the validation set, a key metric for model selection. We employed Weights & Biases (WandB), an advanced tool that supports systematic exploration of parameter space using a Bayesian optimization approach. This strategy enabled us to iteratively test numerous hyperparameter combinations to identify the configurations that delivered optimal performance. Details of these hyperparameter settings with their search ranges and the chosen parameter by the best performing model are provided in the Supplementary Table 2 and 3, respectively (Supplementary Note 3).

Data availability

In accordance with the Health Insurance Portability and Accountability Act (HIPAA), we are unable to distribute the original audio recordings as they might reveal identifiable information about the participants. However, to support transparency and reproducibility, we have made the de-identified extracted features used in this study publicly available at the GitHub repository: <https://github.com/tmadnan/PARK-speech-fusion.git>.

Code availability

The feature extraction pipeline, along with the model training and evaluation scripts used in our study are publicly available at <https://github.com/tmadnan/PARK-speech-fusion.git>.

Received: 15 July 2024; Accepted: 8 April 2025;

Published online: 20 June 2025

References

1. Goetz, C. G. et al. Movement disorder society-sponsored revision of the unified parkinson's disease rating scale (mds-updrs): scale presentation and clinimetric testing results. *Movement disorders: official journal of the Movement Disorder Society* **23**, 2129–2170 (2008).
2. Morel, T. et al. Patient experience in early-stage parkinson's disease: Using a mixed methods analysis to identify which concepts are cardinal for clinical trial outcome assessment. *Neurology and Therapy* **11**, 1319–1340 (2022).
3. Chowdhury, R. N. et al. Pattern of neurological disease seen among patients admitted in tertiary care hospital. *BMC research notes* **7**, 1–5 (2014).
4. Kissani, N. et al. Why does africa have the lowest number of neurologists and how to cover the gap? *Journal of the Neurological Sciences* **434**, 120119 (2022).
5. Dorsey, E. A. et al. Projected number of people with parkinson disease in the most populous nations, 2005 through 2030. *Neurology* **68**, 384–386 (2007).

6. Rahman, W. et al. Detecting parkinson disease using a web-based speech task: Observational study. *Journal of medical Internet research* **23**, e26305 (2021).
7. Dorsey, E. R. et al. Moving parkinson care to the home. *Movement Disorders* **31**, 1258–1262 (2016).
8. Dubey, H., Goldberg, J. C., Abtahi, M., Mahler, L. & Mankodiya, K. Echowear: smartwatch technology for voice and speech treatments of patients with parkinson's disease. In *Proceedings of the conference on Wireless Health*, 1–8 (2015).
9. Yang, Y. et al. Artificial intelligence-enabled detection and assessment of parkinson's disease using nocturnal breathing signals. *Nature medicine* **28**, 2207–2215 (2022).
10. Islam, M. S. et al. Using ai to measure parkinson's disease severity at home. *npj Digital Medicine* **6**, 156 (2023).
11. Schalkamp, A.-K., Peall, K. J., Harrison, N. A. & Sandor, C. Wearable movement-tracking data identify parkinson's disease years before clinical diagnosis. *Nature Medicine* **29**, 2048–2056 (2023).
12. Adnan, T. et al. Unmasking parkinson's disease with smile: An ai-enabled screening framework. *arXiv preprint arXiv:2308.02588* (2023).
13. Wang, X. et al. Alignment by agreement. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, 3184–3192 (2019).
14. Forkan, A. R. M., Branch, P., Jayaraman, P. P. & Ferretto, A. An internet-of-things solution to assist independent living and social connectedness in elderly. *ACM Transactions on Social Computing* **2**, 1–24 (2019).
15. Little, M., Mcsharry, P., Roberts, S., Costello, D. & Moroz, I. Exploiting nonlinear recurrence and fractal scaling properties for voice disorder detection. *Nature Precedings* 1–1 (2007).
16. Saldanha, J. C., Suvarna, M., Satish, D. & Santhmayor, C. Jitter as a quantitative indicator of dysphonia in parkinson's disease. *International Journal of Intelligent Systems Technologies and Applications* **21**, 199–228 (2023).
17. Hawi, S. et al. Automatic parkinson's disease detection based on the combination of long-term acoustic features and mel frequency cepstral coefficients (mfcc). *Biomedical Signal Processing and Control* **78**, 104013 (2022).
18. Upadhy, S. S., Cheeran, A. & Nirmal, J. Statistical comparison of jitter and shimmer voice features for healthy and parkinson affected persons. In *2017 second international conference on electrical, computer and communication technologies (ICECCT)*, 1–6 (IEEE, 2017).
19. Cantürk, İ. & Karabiber, F. A machine learning system for the diagnosis of parkinson's disease from speech signals and its application to multiple speech signal types. *Arabian Journal for Science and Engineering* **41**, 5049–5059 (2016).
20. Rusz, J., Cmejla, R., Ruzickova, H. & Ruzicka, E. Quantitative acoustic measurements for characterization of speech and voice disorders in early untreated parkinson's disease. *The journal of the Acoustical Society of America* **129**, 350–367 (2011).
21. Liu, H., Wang, E. Q., Metman, L. V. & Larson, C. R. Vocal responses to perturbations in voice auditory feedback in individuals with parkinson's disease. *PLoS one* **7**, e33629 (2012).
22. Frid, A., Kantor, A., Svechin, D. & Manevitz, L. M. Diagnosis of parkinson's disease from continuous speech using deep convolutional networks without manual selection of features. In *2016 IEEE International Conference on the Science of Electrical Engineering (ICSEE)*, 1–4 (IEEE, 2016).
23. Umapathy, K., Krishnan, S., Parsa, V. & Jamieson, D. G. Discrimination of pathological voices using a time-frequency approach. *IEEE Transactions on Biomedical Engineering* **52**, 421–430 (2005).
24. Khan, T., Westin, J. & Dougherty, M. Classification of speech intelligibility in parkinson's disease. *Biocybernetics and Biomedical Engineering* **34**, 35–45 (2014).
25. Bai, S., Kolter, J. Z. & Koltun, V. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271* (2018).
26. Ravanelli, M. & Bengio, Y. Speaker recognition from raw waveform with sincnet. In *2018 IEEE spoken language technology workshop (SLT)*, 1021–1028 (IEEE, 2018).
27. Szegedy, C., Ioffe, S., Vanhoucke, V. & Alemi, A. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Proceedings of the AAAI conference on artificial intelligence*, vol. 31 (2017).
28. Ben-Hur, A., Ong, C. S., Sonnenburg, S., Schölkopf, B. & Rätsch, G. Support vector machines and kernels for computational biology. *PLoS computational biology* **4**, e1000173 (2008).
29. Zeng, Z., Pantic, M., Roisman, G. I. & Huang, T. S. A survey of affect recognition methods: audio, visual and spontaneous expressions. In *Proceedings of the 9th international conference on Multimodal interfaces*, 126–133 (2007).
30. Graves, A., Mohamed, A.-r. & Hinton, G. Speech recognition with deep recurrent neural networks. In *2013 IEEE international conference on acoustics, speech and signal processing*, 6645–6649 (IEEE, 2013).
31. Guyon, I. & Elisseeff, A. An introduction to variable and feature selection. *Journal of machine learning research* **3**, 1157–1182 (2003).
32. Paul, S. et al. Bias investigation in artificial intelligence systems for early detection of parkinson's disease: a narrative review. *Diagnostics* **12**, 166 (2022).
33. Rashnu, A. & Salimi-Badr, A. Integrative deep learning framework for parkinson's disease early detection using gait cycle data measured by wearable sensors: A cnn-gru-gnn approach. *arXiv preprint arXiv:2404.15335* (2024).
34. Wang, M. et al. Robust feature engineering for parkinson disease diagnosis: New machine learning techniques. *JMIR Biomedical Engineering* **5**, e13611 (2020).
35. Jhapate, A. K. & Shrivastava, H. Gait based human parkinson's disease detection using fused features with multi-kernel support vector machine. *International Journal of Information Technology* 1–9 (2024).
36. Chen, S. et al. Wavlm: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing* **16**, 1505–1518 (2022).
37. Baevski, A., Zhou, Y., Mohamed, A. & Auli, M. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in neural information processing systems* **33**, 12449–12460 (2020).
38. Girdhar, R. et al. Imagebind: One embedding space to bind them all. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15180–15190 (2023).
39. Hsu, W.-N. et al. Hubert: Self-supervised speech representation learning by masked prediction of hidden units (2021).
40. Klempř, O., Příhoda, D. & Krupička, R. Evaluating the performance of wav2vec embedding for parkinson's disease detection. *Measurement Science Review* **23**, 260–267 (2023).
41. Lin, Y. et al. Efficient transformer for multimodal machine translation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 8066–8076 (2021).
42. Tsai, Y.-H. H., Bai, S., Yamada, M., Morency, L.-P. & Salakhutdinov, R. Transformer dissection: a unified understanding of transformer's attention via the lens of kernel. *arXiv preprint arXiv:1908.11775* (2019).
43. Lu, J., Batra, D., Parikh, D. & Lee, S. VILBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems* **32** (2019).

44. Chen, Y.-C. et al. Uniter: Universal image-text representation learning. *European conference on computer vision* 104–120 (2020).
45. Li, L. H., Yatskar, M., Yin, D., Hsieh, C.-J. & Chang, K.-W. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557* (2019).
46. Langevin, R. et al. The park framework for automated analysis of parkinson's disease characteristics. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* **3**, 1–22 (2019).
47. Vaswani, A. et al. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008 (2017).
48. Wang, X., Girshick, R., Gupta, A. & He, K. Non-local neural networks. *Proceedings of the IEEE conference on computer vision and pattern recognition* 7794–7803 (2018).
49. Maćkiewicz, A. & Ratajczak, W. Principal components analysis (PCA). *Computers & Geosciences* **19**, 303–342 (1993).
50. Adnan, T. M. T., Tanjim, M. M. & Adnan, M. A. Fast, scalable and geo-distributed PCA for big data analytics. *Information Systems* **98**, 101710 (2021).
51. Chawla, N. V., Bowyer, K. W., Hall, L. O. & Kegelmeyer, W. P. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research* **16**, 321–357 (2002).
52. Mohammed, R., Rawashdeh, J. & Abdullah, M. Machine learning with oversampling and undersampling techniques: overview study and experimental results. In *2020 11th international conference on information and communication systems (ICICS)*, 243–248 (IEEE, 2020).
53. Armstrong, R. A. When to use the bonferroni correction. *Ophthalmic and Physiological Optics* **34**, 502–508 (2014).
54. Upton, G. J. Fisher's exact test. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **155**, 395–402 (1992).
55. Kwak, S. G. & Kim, J. H. Central limit theorem: the cornerstone of modern statistics. *Korean journal of anesthesiology* **70**, 144–156 (2017).
56. Shiryayev, A., Aronowitz, H., Walach, E., Chowdhury, S. & Sharma, A. Responsible machine learning with error analysis. (2021). Accessed: 2024-05-10.
57. Taylor, P. Number of smartphone mobile network subscriptions worldwide from 2016 to 2022, with forecasts from 2023 to 2028. *Temmuz* **30**, 2023 (2023).
58. Yang, W. et al. Current and projected future economic burden of parkinson's disease in the us. *npj Parkinson's Disease* **6**, 15 (2020).
59. Fujita, T., Babazono, A., Kim, S., Jamal, A. & Li, Y. Effects of physician visit frequency for parkinson's disease treatment on mortality, hospitalization, and costs: a retrospective cohort study. *BMC geriatrics* **21**, 1–12 (2021).
60. Dashtipour, K., Tafreshi, A., Lee, J. & Crawley, B. Speech disorders in parkinson's disease: pathophysiology, medical management and surgical approaches. *Neurodegenerative disease management* **8**, 337–348 (2018).
61. Rizzo, G. et al. Accuracy of clinical diagnosis of parkinson disease: a systematic review and meta-analysis. *Neurology* **86**, 566–576 (2016).
62. Bayram, F., Ahmed, B. S. & Kassler, A. From concept drift to model degradation: An overview on performance-aware drift detectors. *Knowledge-Based Systems* **245**, 108632 (2022).
63. Ackerman, S., Dube, P., Farchi, E., Raz, O. & Zalmanovici, M. Machine learning model drift detection via weak data slices. In *2021 IEEE/ACM Third International Workshop on Deep Learning for Testing and Testing for Deep Learning (DeepTest)*, 1–8 (IEEE, 2021).
64. Hadsell, R., Rao, D., Rusu, A. A. & Pascanu, R. Embracing change: Continual learning in deep neural networks. *Trends in cognitive sciences* **24**, 1028–1040 (2020).
65. Aleksovski, D., Miljkovic, D., Bravi, D. & Antonini, A. Disease progression in parkinson subtypes: the ppmi dataset. *Neurological Sciences* **39**, 1971–1976 (2018).
66. Chatterjee, I. & Kumar, S. An analytical study of age and gender effects on voice range profile in bengali adult speakers using phonetogram. *International Journal of Phonosurgery and Laryngology* (2011).
67. Taylor, S., Dromey, C., Nissen, S. & Tanner, K. Age-related changes in speech and voice: Spectral and cepstral measures. *Journal of speech, language, and hearing research : JSLHR* **63**, 647–660 (2020).
68. Stathopoulos, E. T., Huber, J. E. & Sussman, J. E. Age-related changes in speech and voice: Spectral and cepstral measures. *Journal of speech, language, and hearing research : JSLHR* **54**, 1011–1021 (2011).
69. Blagus, R. & Lusa, L. Smote for high-dimensional class-imbalanced data. *BMC bioinformatics* **14**, 1–16 (2013).
70. Lundberg, S. & Lee S. A unified approach to interpreting model predictions. *Advances in neural information processing systems* **30** (2017).
71. Ribeiro, M. T., Singh, S. & Guestrin, C. " why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 1135–1144 (2016).
72. Gohel, P., Singh, P. & Mohanty, M. Explainable ai: current status and future directions. *arXiv preprint arXiv:2107.07045* (2021).
73. Li, P., Yang, Y., Pagnucco, M. & Song, Y. Explainability in graph neural networks: An experimental survey. *arXiv preprint arXiv:2203.09258* (2022).
74. Kim, J., Kumar, M., Gowda, D., Garg, A. & Kim, C. Semi-supervised transfer learning for language expansion of end-to-end speech recognition models to low-resource languages. In *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 984–988 (IEEE, 2021).
75. Torre, I. G., Romero, M. & Álvarez, A. Improving aphasic speech recognition by using novel semi-supervised learning methods on aphasiabank for english and spanish. *Applied Sciences* **11**, 8872 (2021).
76. Frost, G., Morris, E., Jansen van Vuren, J. & Niesler, T. Fine-tuned self-supervised speech representations for language diarization in multilingual code-switched speech. In *Southern African Conference for Artificial Intelligence Research*, 246–259 (Springer, 2022).
77. Hallett, M. & Khoshbin, S. A physiological mechanism of bradykinesia. *Brain* **103**, 301–314 (1980).
78. Maycas-Cepeda, T. et al. Hypomimia in parkinson's disease: what is it telling us? *Frontiers in Neurology* **11**, 603582 (2021).
79. Radford, A. et al. Robust speech recognition via large-scale weak supervision. In *International Conference on Machine Learning*, 28492–28518 (PMLR, 2023).
80. Rabiner, L. Digital processing of speech signals. *Prentice Hall google schola* **2**, 601–604 (1978).
81. Sataloff, R. T. *Voice science* (Plural Publishing, 2017).
82. Titze, I. R. & Martin, D. W. Principles of voice production (1998).
83. Boersma, P. & Van Heuven, V. Speak and unspeak with praat. *Glott International* **5**, 341–347 (2001).
84. Styler, W. Using praat for linguistic research. *University of Colorado at Boulder Phonetics Lab* (2013).
85. Jadoul, Y., Thompson, B. & De Boer, B. Introducing parselmouth: A python interface to praat. *Journal of Phonetics* **71**, 1–15 (2018).
86. Chen, C. et al. Unispeech: Unified speech representation learning with labeled and unlabeled data. In *International Conference on Machine Learning*, 1719–1729 (PMLR, 2021).
87. Babu, A. et al. Xls-r: Self-supervised cross-lingual speech representation learning at scale. In *Proceedings of the IEEE/CVF*

- Conference on Computer Vision and Pattern Recognition, 17197–17206 (2021).
88. Cohen, I., Huang, Y., Chen, J. & Benesty, J. Pearson correlation coefficient. In *Noise Reduction in Speech Processing*, 1–4 (Springer, 2009).
 89. Jain, A. K. Data clustering: 50 years beyond k-means. *Pattern Recognition Letters* **31**, 651–666 (2010).
 90. Patro, S. Normalization: A preprocessing stage. *arXiv preprint arXiv:1503.06462* (2015).
 91. Raschka, S. & Mirjalili, V. *Python machine learning: Machine learning and deep learning with Python, scikit-learn, and TensorFlow 2* (Packt publishing ltd, 2019).
 92. Han, J., Pei, J. & Tong, H. *Data mining: concepts and techniques* (Morgan kaufmann, 2022).
 93. Scholkopf, B. & Smola, A. *J. Learning with kernels: support vector machines, regularization, optimization, and beyond* (MIT press, 2018).
 94. LeCun, Y., Bottou, L., Orr, G. B. & Müller, K.-R. Efficient backprop. In *Neural networks: Tricks of the trade*, 9–50 (Springer, 2002).
 95. Ngiam, J. et al. Multimodal deep learning. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, 689–696 (2011).
 96. Baltrušaitis, T., Ahuja, C. & Morency, L.-P. Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence* **41**, 423–443 (2019).
 97. Hinton, G. E. & Salakhutdinov, R. R. Reducing the dimensionality of data with neural networks. *Science* **313**, 504–507 (2006).
 98. Zadeh, A., Chen, M., Poria, S., Cambria, E. & Morency, L.-P. Tensor fusion network for multimodal sentiment analysis. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 1103–1114 (2017).
 99. Tsai, Y.-H. H., Liang, P. P., Zadeh, A., Salakhutdinov, R. & Morency, L.-P. Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 6984–6991 (2019).
 100. Lei, J. et al. Univl: A unified video and language pre-training model for multimodal understanding and generation. *arXiv preprint arXiv:2202.03828* (2022).
 101. Tan, H. & Bansal, M. Lxmert: Learning cross-modality encoder representations from transformers. *arXiv preprint arXiv:1908.07490* (2019).
 102. Chen, J. & Ho, C. M. Mm-vit: Multi-modal video transformer for compressed video action recognition. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, 1910–1921 (2022).
 103. Geng, X. et al. Multimodal masked autoencoders learn transferable representations. *arXiv preprint arXiv:2205.14204* (2022).
 104. Bergstra, J. & Bengio, Y. Random search for hyper-parameter optimization. *Journal of machine learning research* **13** (2012).
 105. Hlavnička, J. et al. Automated analysis of connected speech reveals early biomarkers of parkinson's disease in patients with rapid eye movement sleep behaviour disorder. *Scientific reports* **7**, 1–13 (2017).
 106. De Boer, P.-T., Kroese, D. P., Mannor, S. & Rubinstein, R. Y. A tutorial on the cross-entropy method. *Annals of operations research* **134**, 19–67 (2005).
 107. Wang, H. et al. Cosface: Large margin cosine loss for deep face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5265–5274 (2018).
 108. Goodfellow, I., Bengio, Y. & Courville, A. *Deep learning* (MIT press, 2016).
 109. Kendall, A., Gal, Y. & Cipolla, R. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7482–7491 (2018).
 110. Chen, R. et al. Imagebind: One embedding space to bind them all. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 15180–15190 (2023).
 111. Yosinski, J., Clune, J., Bengio, Y. & Lipson, H. How transferable are features in deep neural networks? In *Advances in neural information processing systems*, 3320–3328 (2014).
 112. Hawkins, D. M. The problem of overfitting. *Journal of chemical information and computer sciences* **44**, 1–12 (2004).
 113. Larochelle, H., Bengio, Y., Louradour, J. & Lamblin, P. Exploring strategies for training deep neural networks. In *Journal of machine learning research* (2009).
 114. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I. & Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* **15**, 1929–1958 (2014).
 115. Sze, V., Chen, Y.-H., Yang, T.-J. & Emer, J. S. Efficient processing of deep neural networks: A tutorial and survey. *Proceedings of the IEEE* **105**, 2295–2329 (2017).
 116. Zhang, D., Guo, Y., Cheng, X., Cao, C. & Fang, Y. Real-time digital signal processing based on fpgas for electronic skin implementation. *Sensors* **18**, 1635 (2018).
 117. Diamant, N. et al. Patient contrastive learning: A performant, expressive, and practical approach to electrocardiogram modeling. *PLoS computational biology* **18**, e1009862 (2022).
 118. Liu, H., Li, C., Wu, Q. & Lee, Y. J. Visual instruction tuning. *Advances in neural information processing systems* **36**, 34892–34916 (2023).
 119. Victoria, A. H. & Maragatham, G. Automatic tuning of hyperparameters using bayesian optimization. *Evolving Systems* **12**, 217–223 (2021).
 120. Feurer, M. & Hutter, F. Hyperparameter optimization. *Automated machine learning: Methods, systems, challenges* 3–33 (2019).

Acknowledgements

We express our profound gratitude to Meghan Pawlik for her substantial contributions to the maintenance of the clinical studies and the data collection. We are also deeply grateful to Dr. Ruth Schneider, Dr. Ray Dorsey, and Emily Hartman from the University of Rochester Medical Center (URMC) for their invaluable support in facilitating the collection of H&Y PD stage information from the clinical records of study participants. Similarly, we extend our appreciation to Cathie Schwartz and Karen Jaffe from the designated PD Care Facility, whose efforts were instrumental in gathering both participant data and PD stage information. Additionally, we are thankful to Sangwu Lee for his meticulous feedback on the manuscript and his invaluable assistance in debugging the code. The project was supported by the National Institute of Neurological Disorders and Stroke of the National Institutes of Health under award number P50NS108676, Gordon and Betty Moore Foundation, and Google Faculty Research Award. One of the authors is supported by a Google PhD Fellowship.

Author contributions

TA, AA, and MSI conceptualized the study and designed the experimental framework. AA and MSI were responsible for developing the feature extraction pipelines, while also collaborating on the development of the baseline machine learning models using individual feature sets. AA crafted the fusion models using concatenation techniques, whereas TA and EH' (Ekram Hossain) focused on advancing the projection-based fusion architecture. Together, AA, TA, and EH' engaged in model evaluation, results analysis, visualization, and manuscript preparation. RL and SP significantly contributed to error visualization and statistical analysis. All authors critically reviewed the manuscript, providing substantial revisions and improvements as needed. EH (Ehsan Hoque), as the principal investigator, along with MSI,

supervised the entire project, providing strategic direction and refining the manuscript's narrative.

Competing interests

All the authors declare that there are no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41531-025-00956-7>.

Correspondence and requests for materials should be addressed to Ehsan Hoque.

Reprints and permissions information is available at <http://www.nature.com/reprints> 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-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, 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 you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. 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-nc-nd/4.0/>.

© The Author(s) 2025