

Accessible, At-Home Detection of Parkinson’s Disease via Multi-Task Video Analysis

Md Saiful Islam^{1*}, Tariq Adnan¹, Jan Freyberg², Sangwu Lee¹, Abdelrahman Abdelkader¹, Meghan Pawlik³, Cathe Schwartz⁴, Karen Jaffe⁴, Ruth B. Schneider³, Ray Dorsey³, Ehsan Hoque¹

¹University of Rochester, New York, United States

²Google Research, London, United Kingdom

³ University of Rochester Medical Center, New York, United States

⁴ InMotion, Ohio, United States

mislam6@ur.rochester.edu

Abstract

Limited accessibility to neurological care leads to under-diagnosed Parkinson’s Disease (PD), preventing early intervention. Existing AI-based PD detection methods primarily focus on unimodal analysis of motor or speech tasks, overlooking the multifaceted nature of the disease. To address this, we introduce a large-scale, multi-task video dataset of 1102 sessions (each containing videos of finger tapping, facial expression, and speech tasks captured via webcam) from 845 participants (272 with PD). We propose a novel Uncertainty-calibrated Fusion Network (*UFNet*) that leverages this multimodal data to enhance diagnostic accuracy. *UFNet* employs independent task-specific networks, trained with Monte Carlo Dropout for uncertainty quantification, followed by self-attended fusion of features, with attention weights dynamically adjusted based on task-specific uncertainties. We randomly split the participants into training (60%), validation (20%), and test (20%) sets to ensure patient-centered evaluation. *UFNet* significantly outperformed single-task models in terms of accuracy, area under the ROC curve (AUROC), and sensitivity while maintaining non-inferior specificity. Withholding uncertain predictions further boosted the performance, achieving $88.0 \pm 0.3\%$ accuracy, $93.0 \pm 0.2\%$ AUROC, $79.3 \pm 0.9\%$ sensitivity, and $92.6 \pm 0.3\%$ specificity, at the expense of not being able to predict for $2.3 \pm 0.3\%$ data (\pm denotes 95% confidence interval). Further analysis suggests that the trained model does not exhibit any detectable bias across sex and ethnic subgroups and is most effective for individuals aged between 50 and 80. By merely requiring a webcam and microphone, our approach facilitates accessible home-based PD screening, especially in regions with limited healthcare resources.

Code & Dataset — <https://github.com/ROC-HCI/UFNet>

Extended version — <https://arxiv.org/abs/2406.14856>

Demo — <https://parktest.net/demo>

Introduction

Due to limited access to neurological care, many individuals, particularly in underserved regions (Kissani et al. 2022), live

*Corresponding author

with Parkinson’s disease—the fastest growing neurological disorder (Dorsey et al. 2018)—without even knowing it. In many instances, when diagnosed late with this incurable disease, the condition could have already progressed significantly, limiting the usefulness of available medications. Imagine a future where individuals can remotely assess their risk for PD by simply visiting a website, activating their webcam and microphone, and completing a series of standardized tasks. This accessible approach could empower people to seek early intervention and treatment, potentially improving their quality of life.

Detecting PD is particularly challenging due to the variability of individual symptoms. For example, while vocal impairment is a common PD symptom (Ho et al. 1999), patients may also exhibit PD through other modalities, such as facial expression (e.g., hypomimia (Gunnery et al. 2016)) or motor function (e.g., bradykinesia (Bologna et al. 2023)). Consequently, models relying solely on a single task may yield suboptimal performance. To address this, we introduce a large video dataset featuring webcam recordings of individuals performing three tasks: (i) finger-tapping (motor function), (ii) smiling (facial expression), and (iii) uttering a pangram¹ (speech). The dataset is collected from 845 unique participants (272 with PD) from diverse demographics who recorded all of these tasks successfully (about 20% of them recorded multiple times), resulting in 1102 videos for each task and a total of 3306 videos. To our knowledge, this is the first multi-task video dataset for PD screening.

We propose a novel two-stage classification model to distinguish between individuals with and without Parkinson’s disease (PD) using our dataset. First, each task is independently modeled with neural networks trained using Monte Carlo dropout (MC dropout) to generate predictions and associated uncertainties. Next, we introduce the Uncertainty-calibrated Fusion Network (*UFNet*), which aggregates features from multiple tasks through an attention mechanism (Vaswani et al. 2017), while calibrating the attention scores based on task-specific uncertainties. *UFNet* is designed to produce PD/Non-PD predictions and uses MC dropout during training and

¹pangram: sentence containing all the letters of the alphabet.

inference to quantify prediction confidence. By filtering out low-confidence predictions, it increases patient safety. Our evaluations on a subject-separated test set demonstrate that *UFNet* significantly outperforms single-task models and multi-modal fusion baselines. The model is computationally efficient, employing only a few shallow neural networks and a single self-attention module, making it suitable for deployment on smartphones or personal computers.

Here is a summary of our key contributions:

- We introduce the first-ever large-scale, multi-task video dataset for Parkinson’s disease screening collected from 845 individuals with diverse demography. Although we cannot share the raw videos to protect identifiable patient information, we release the dataset as extracted features.
- We propose a novel, effective multimodal fusion model named *UFNet*, achieving $88.0 \pm 0.3\%$ accuracy and $93.0 \pm 0.2\%$ AUROC on a subject-separated test dataset. The proposed model significantly outperformed single-task models and multimodal fusion baselines.

Related Works

Traditionally, Parkinson’s Disease (PD) is diagnosed by a clinician based on medical history and a clinical examination, typically including completing standardized tasks and rating each task according to the MDS-UPDRS guidance (Goetz et al. 2008). Recently, a cerebrospinal fluid (CSF) based α -synuclein seed amplification assay has been developed (Siderowf et al. 2023), offering a potential diagnostic biomarker. However, access to clinical care can be limited, and diagnostic methods relying on the collection of CSF are invasive, costly, and burdensome to the patients.

Recent research has leveraged machine learning (ML) and sensors to enable remote assessment of PD. For example, breathing signals obtained from reflected radio waves can detect PD accurately when analyzed by an ML model (Yang et al. 2022). Additionally, body-worn sensors have been successfully used to monitor clinical features such as dyskinesia and gait disturbances associated with PD (Moreau et al. 2023). However, wearable sensors face challenges related to cost, comfort, and ease of use, which limit the scalability of these techniques for global adoption.

Advanced pose tracking models like MediaPipe (Lugaresi et al. 2019) enable extracting precise clinical features from recorded videos, later used for PD screening (Jin et al. 2020) or progression tracking (Islam et al. 2023). However, existing methods for video analysis suffer from major limitations – small cohort size (Jin et al. 2020), and reliance on single modality (Tsanas et al. 2009; Ali et al. 2020, 2021; Rahman et al. 2021; Adnan et al. 2023). Symptoms of PD are multi-faceted and may affect individuals differently. For instance, one individual may face speech difficulty but retain normal motor functionality, while others may have prominent hypomimia (i.e., reduced facial expression) or bradykinesia (i.e., slowness of movement). Therefore, PD detection models may need to consider multiple modalities for improved efficacy.

This work addresses a critical gap in the field, which has persisted due to the lack of multimodal, diverse, and

naturalistic data (Sibley et al. 2021). By utilizing the PARK framework (Langevin et al. 2019) and inter-disciplinary collaboration, we have collected videos of multiple tasks from a large, diverse cohort, enabling us to develop an accurate PD screening model. We hope the community will engage with this dataset and continue to advance the model.

Dataset

Standardized Tasks

We select three standardized tasks that can be easily completed using a computer webcam and microphone, with or without external supervision:

Finger-tapping: Participants tap their thumb with the index finger ten times as fast as possible, first with the right hand, then with the left. The finger-tapping task is completed following the MDS-UPDRS scale to measure bradykinesia in the upper limb, a key sign of PD (Hughes et al. 1992).

Smile: Participants mimic a smile expression three times, alternating with a neutral face. Studies suggest the task captures signs of hypomimia (Bandini et al. 2017; Adnan et al. 2023), even though the expression may be unnatural.

Speech: Participants utter a script, “*The quick brown fox jumps over a lazy dog. The dog wakes up and follows the fox into the forest. But again, the quick brown fox jumps over the lazy dog.*” The first sentence is a pangram, containing all the letters of the English alphabet, and the later sentences are added to obtain a longer speech segment. Prior research identified this task as a promising way of screening PD (Rahman et al. 2021).

Participant Recruitment

We recruited participants (with and without PD) from a diverse background through multiple channels, including a brain health study registry managed by the University of Rochester Medical Center (URMC), In-Motion (a PD wellness center in Ohio), clinician referrals, and social media outreach. Among the approximately 1,400 individuals who recorded at least one task, 845 completed all three standardized tasks (comprising this dataset). Participants are primarily from the US, with diverse demographic characteristics: 52.7% female, a mean age of 61.9 years (22.5% over 70), and representation from various ethnic backgrounds. Among the 272 (32.2%) participants who had PD, 233 had their diagnosis confirmed through clinical evaluation, while 39 self-reported their condition. This study was conducted in collaboration with URMC and was approved by the institutional review board (IRB) of the University of Rochester.

Dataset Splits

Most participants completed all three tasks, but some did not, and some task videos were discarded due to feature extraction failures, often caused by inaccurate task completion. Instead of discarding videos of participants with missing tasks, we trained task-specific models with all available videos for each task. The multi-task model is then trained on the participants who completed all three tasks (Table 1). We split the datasets based on the participants

Task/Split	#Sessions, PD %	#Participants, PD %
Finger-tapping	1374, 41.3%	1167, 36.6%
Training	945, 43.9%	819, 38.8%
Validation	221, 37.6%	172, 32.6%
Test	208, 33.2%	176, 30.1%
Smile	1684, 32.8%	1357, 28.5%
Training	1021, 33.2%	824, 28.4%
Validation	342, 33.9%	266, 28.6%
Test	321, 30.5%	267, 28.8%
Speech	1655, 33.9%	1265, 28.9%
Training	1007, 35.3%	769, 29.0%
Validation	338, 33.7%	252, 29.0%
Test	310, 29.7%	244, 28.7%
All tasks	1102, 41.8%	845, 32.2%
Training	690, 45.1%	516, 32.6%
Validation	215, 38.1%	167, 32.9%
Test	197, 34.9%	162, 30.2%

Table 1: Summary of dataset splits.

to ensure patient-centric evaluation. All participants were randomly assigned into training (60%), validation (20%), and test (20%) sets, with stratification to maintain a similar ratio of individuals with and without PD across folds. Both task-specific and multi-task models are validated and tested on the same participant cohort, with no model seeing any data from these participants during training.

Our Approach

Feature Extraction

We rely on prior literature to extract task-specific video features. For the smile and finger-tapping tasks, clinically meaningful hand-crafted features were the better choice against deep video models such as ViViT (Arnab et al. 2021), TimeSformer (Bertasius, Wang, and Torresani 2021), VideoMAE (Tong et al. 2022), and Uniformer (Li et al. 2023), based on their performance and suitability for low-resource deployment. WavLM (Chen et al. 2022) has been identified as the most effective for speech-based PD screening by a recent work (Adnan et al. 2024).

Finger-tapping features. Islam et al. (2023) extracted 65 features to analyze the finger-tapping task for assessing PD severity. Using MediaPipe hand (Grishchenko et al. 2022) to detect hand movements and key points, they measured clinically relevant features such as finger-tapping speed, amplitude, and interruptions. We apply this technique to both hands, extracting 130 features in total.

Smile features. We used 42 facial features from smile videos extracted by Adnan et al. (2023). These features, extracted using OpenFace (Baltrušaitis, Robinson, and Morency 2016) and MediaPipe, capture key PD markers such as eye blinking, lip separation, mouth opening, and intensity of facial muscle movements.

Speech features. We extracted 1024-dimensional embeddings from a pre-trained WavLM (Chen et al. 2022) model to encode the speech task. WavLM, trained on massive amounts of speech data, excels at capturing

acoustic characteristics of speech, making it useful for tasks like speaker identification, speech, and emotion recognition.

Task-Specific Models

Each task utilizes a separate machine learning model to distinguish between individuals with and without PD. These models have three main components:

- **Feature selection and scaling:** We calculate pairwise correlation among the features to drop highly correlated features. Feature values are then scaled using *StandardScaler* or *MinMaxScaler*. Whether to apply feature selection or scaling, the correlation threshold and the scaling method are hyper-parameters tuned on the validation set.
- **Shallow neural networks:** The network consists of 1-2 linear layers, with a single output neuron in the last layer. We use `sigmoid` activation with the output layer and `ReLU` with the other layers to add non-linearity.
- **Monte Carlo dropout (MC dropout):** To improve model robustness and estimate prediction uncertainty, we employ MC dropout (Atighehchian et al. 2022). This technique allows using dropout for both training and inference. Therefore, we can obtain predictions from multiple rounds during testing and the mean (μ) and standard deviation (σ) of these predictions. We expect a relatively higher standard deviation across these predictions when the model is uncertain. Therefore, σ is used to estimate model uncertainty.

The task-specific models are trained using binary cross-entropy loss and either SGD or AdamW optimizer.

Uncertainty-Calibrated Fusion Network, *UFNet*

The fully trained task-specific models remain frozen during the training of *UFNet*. For each task i , the extracted features ($X_i \in \mathbb{R}^{d_{x_i}}$), predicted mean probability (μ_i), and uncertainty in the prediction (σ_i) are input to *UFNet*. The model then combines information from all the tasks through a series of steps to generate a final, more robust detection of PD (see Figure 1 for an overview).

Projection. Since the size of the features may vary from task to task, they are first projected to the same dimension (d) using a projection layer. Each projection layer consists of a linear layer ($\mathbb{R}^{d_{x_i}} \rightarrow \mathbb{R}^d$) with MC dropout, followed by non-linear activation (`ReLU`) and layer normalization.

Calibrated self-attention. We employ self-attention to integrate the projected features ($X_i^P \in \mathbb{R}^d$) extracted from three distinct tasks. Unlike standard self-attention applications, our input sequence comprises task-specific feature vectors (sequence length = 3) rather than tokens of the same modality. For each task, projected features are transformed into query (Q), key (K), and value (V) vectors using linear projections. Attention weights are computed based on the similarity between queries and keys.

To prioritize informative tasks, we adjust the attention scores to down-weight contributions from tasks with higher prediction uncertainty. Specifically, in addition to using

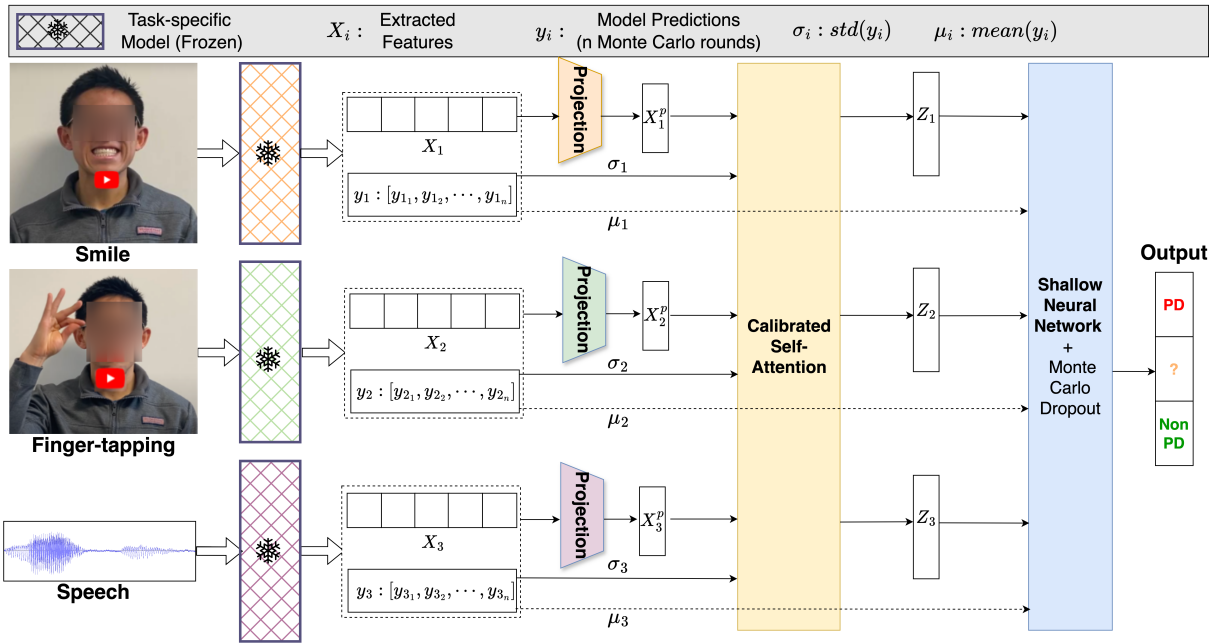


Figure 1: **An illustrative overview of the proposed model.** Task-specific models are shallow neural networks trained with task-specific features and MC dropout. Standard deviations in multiple rounds of inference obtained from the task-specific models are used to calibrate the attention scores when fusing task-specific features. Finally, another shallow neural network is trained to differentiate individuals with and without PD and withhold prediction when the model is uncertain.

traditional query-key similarity, we update the attention matrix A with task-specific uncertainty scores $\Sigma = [\sigma_1, \sigma_2, \sigma_3]$ as,

$$A = \text{softmax}\left(\frac{Q \cdot K^T}{\sqrt{d}} - \eta \Sigma\right)$$

Hyper-parameter η controls the degree of uncertainty-based down-weighting. After computing the attention scores, X_i^p is converted into a contextualized representation $Z_i = \sum_{j=1}^N A_{i,j} V_j$, where $N = 3$ is the number of tasks.

Shallow neural network. The contextualized representations (Z_1, Z_2, Z_3) obtained after self-attention are concatenated along with the task-specific (mean) predicted probabilities (μ_1, μ_2, μ_3). The merged vector is then input to a shallow neural network similar to the one used for task-specific training. The network is trained with 30 rounds of MC dropout, and the average output is used as the final prediction (PD if the average output is more than 0.50, non-PD otherwise).

Withholding predictions. Our proposed approach is designed for neurological screening, specifically to suggest clinical evaluation for individuals if the model identifies PD symptoms. Since mispredictions in this context could lead to significant consequences (e.g., unnecessary anxiety or delayed care), the model is designed to withhold predictions it identifies as uncertain. To achieve this, we implement two different techniques:

- **MC Dropout:** For new test data, we generate n predictions by running multiple rounds of Monte Carlo

(MC) dropout. We compute the 95% confidence interval (CI) for these predictions and classify them as uncertain if the CI contains the decision threshold (i.e., 0.50).

- **Conformal Prediction:** We partition the training data into a training set (80%) and a calibration set (20%). We calculate a conformity threshold using the calibration set and identify uncertain predictions when the prediction set contains both PD and Non-PD classes.

We also experiment with performing label smoothing (Müller, Kornblith, and Hinton 2019) (which redistributes probability mass to reduce overconfidence) and/or Platt scaling (Platt et al. 1999) (which fits a logistic regression model to improve calibration) with conformal prediction (Saunders, Gammernan, and Vovk 1999) to see whether they further improve model performance.

Hyper-parameter tuning. We utilized Weights & Biases² to employ Bayesian hyper-parameter tuning within the defined search space to maximize the AUROC on the validation set (see extended version). Finally, the *UFNet* model is trained with binary cross-entropy loss and SGD optimizer with learning rate 0.0207 and momentum 0.6898. We selected 64 as the dimension for the query, key, and value vectors. The chosen value of dropout probability was 0.4960, and η was 81.8179. Since some hyper-parameters had continuous search space (e.g., the learning rate was tuned within $1e^{-5}$ and 1.0), we report the precise values of the selected hyper-parameters.

²<https://wandb.ai>

Experiment Setup

Experiments are conducted on an AMD Threadripper 3970x 32-core CPU with 256 GB RAM and two NVIDIA A6000 GPUs with 48 GB VRAM. We use Python (PyTorch deep learning framework) and Linux operating system.

Multimodal Baselines

We compare our proposed *UFNet* with four popular multimodal fusion approaches.

Majority Voting. The three task-specific models’ predictions are combined to generate a single prediction. The final prediction is the class (PD or Non-PD) agreed upon by the majority (i.e., two or more).

Neural Late Fusion. The logit scores from the task-specific models are input to a shallow neural network trained to predict a binary class, similar to logistic regression or other ensembling methods.

Early Fusion Baseline. Features from all three tasks are concatenated and input to a shallow neural network.

Hybrid Fusion Baseline. Task-specific features and prediction scores (logits) are provided as input to a shallow neural network. This network leverages both the input features and prediction scores, combining the strengths of early and late fusion approaches.

Performance Reporting

After hyper-parameter tuning, we train the selected model with 30 random seeds and evaluate it on the test set. We report the average and 95% confidence interval (CI) of each performance metric. In addition to standard binary classification metrics, we report coverage (% of cases where the model is confident enough to make a prediction) when appropriate. Furthermore, ensuring that the class probabilities output by a model approximates the actual probability is critical for patient safety. Therefore, we evaluate model calibration using expected calibration error (ECE) (Nixon et al. 2019) and Brier score (Rufibach 2010). Unless specified otherwise, a significant difference is identified by having no overlap in the 95% CI, while non-inferiority indicates overlap.

Results

Task-Specific Model Performance

Among the three standardized tasks, speech is the most accurate for PD classification, and finger-tapping is the least accurate. The shallow neural network achieved $84.5 \pm 0.3\%$ accuracy and $89.4 \pm 0.2\%$ AUROC using only speech features. Models trained on right-hand tapping features performed better than those trained on left-hand features, indicating potential under-representation of left-handed individuals. However, the model trained on concatenated features from both hands significantly improved the F_1 score while being non-inferior in other metrics (see extended version). The finger-tapping model refers to the one trained on both-hand features for the remaining analysis.

MC dropout significantly boosted the performance of the smile model in all metrics. For the speech model, while MC dropout improved accuracy and balanced accuracy,

Task	Accuracy	F_1 score	AUROC
Finger-tapping			
w/o MC dropout	72.0 ± 0.9	60.2 ± 1.3	73.9 ± 0.9
w/ MC dropout	73.1 ± 0.7	61.7 ± 0.9	74.9 ± 0.7
Smile			
w/o MC dropout	75.6 ± 0.2	64.3 ± 0.3	83.2 ± 0.1
w/ MC dropout	<u>77.6 ± 0.2</u>	<u>67.5 ± 0.3</u>	<u>83.6 ± 0.1</u>
Speech			
w/o MC dropout	84.5 ± 0.3	71.7 ± 0.4	89.4 ± 0.2
w/ MC dropout	85.1 ± 0.2	72.1 ± 0.6	87.8 ± 0.1

Table 2: Performance of task-specific models with (w/) and without (w/o) MC dropout (in percentage). Underlined metrics denote significantly better performance (within the same task), while **bold** metrics show overall best performance.

Task Combination	Accuracy	F_1 score	AUROC
All three tasks	87.3 ± 0.4	81.0 ± 0.6	92.8 ± 0.2
Tapping + Smile	78.0 ± 0.8	65.6 ± 1.7	84.8 ± 0.5
Tapping + Speech	84.1 ± 0.3	77.3 ± 0.4	91.4 ± 0.2
Smile + Speech	85.2 ± 0.3	75.0 ± 0.4	91.2 ± 0.1

Table 3: Performance of models (as percentages) trained on different combinations of the standardized tasks. “Tapping” is used as the short form of the finger-tapping task. The **bold** metrics denote best performance.

it decreased AUROC and AUPRC. The performance of the finger-tapping model was not notably affected. Due to the additional benefits of MC dropout in modeling prediction uncertainty, the multi-task models (*UFNet* and other baselines) were trained with MC dropout unless specified otherwise. See Table 2 for details.

Effect of Multi-task Combinations

Combining multiple tasks using the proposed *UFNet* model enhanced performance (Table 3). For instance, the AUROC scores for all multi-task models were significantly better than the corresponding single-task scores. Although the finger-tapping task alone was the weakest for detecting PD, its features may have complemented other task features, resulting in significant improvements in most metrics. Combining all three tasks significantly improved all reported metrics, achieving an accuracy of $87.3 \pm 0.3\%$, AUROC of $92.8 \pm 0.2\%$, and an F_1 score of $81.0 \pm 0.6\%$.

Comparison Against Baselines

The proposed uncertainty-calibrated fusion network (*UFNet*) outperformed the baseline methods in most metrics. More specifically, *UFNet* significantly improved PD classification accuracy, precision, and specificity compared to any of the four baselines. It also achieved higher balanced accuracy, AUROC, and F_1 score while maintaining statistically non-inferior sensitivity (Table 4).

Ablation Results

To investigate the significance of our introduced uncertainty-calibrated self-attention module, we replace

Model	Accuracy	Balanced Accuracy	AUROC	F ₁ score	Precision	Sensitivity (Recall)	Specificity
Baseline models							
Majority Voting	85.3	83.9	89.6	78.2	80.0	76.5	89.9
Neural Late Fusion	84.1 ± 0.4	81.3 ± 4.8	91.7 ± 2.2	73.2 ± 8.3	73.5 ± 7.5	76.3 ± 9.4	88.2 ± 2.7
Early Fusion Baseline	83.6 ± 0.6	81.8 ± 0.7	91.0 ± 0.2	76.7 ± 0.7	75.4 ± 1.1	78.1 ± 0.9	86.5 ± 0.8
Hybrid Fusion Baseline	84.1 ± 0.3	82.4 ± 0.4	91.4 ± 0.2	77.3 ± 0.4	76.2 ± 0.7	78.6 ± 0.6	87.0 ± 0.6
Our model with different attention modules							
Dot product self-attention (Vaswani et al. 2017)	85.5 ± 0.4	84.3 ± 0.4	92.9 ± 0.2	78.3 ± 0.6	80.7 ± 0.6	76.1 ± 1.1	90.4 ± 0.4
LRFormer (Ye et al. 2023)	86.2 ± 0.5	85.1 ± 0.5	92.6 ± 0.3	79.5 ± 0.7	81.7 ± 0.9	77.6 ± 1.0	90.8 ± 0.6
UFNet (ours)	87.3 ± 0.4	86.4 ± 0.4	92.8 ± 0.2	81.0 ± 0.6	83.8 ± 0.5	78.4 ± 1.0	92.0 ± 0.3
Early vs. hybrid fusion in UFNet							
Early Fusion	86.7 ± 0.5	85.8 ± 0.5	92.7 ± 0.3	79.9 ± 0.8	83.3 ± 0.7	76.9 ± 1.4	91.9 ± 0.4
Hybrid Fusion	87.3 ± 0.4	86.4 ± 0.4	92.8 ± 0.2	81.0 ± 0.6	83.8 ± 0.5	78.4 ± 1.0	92.0 ± 0.3

Table 4: Comparison of *UFNet* performance (as percentages) against multimodal baselines and other attention modules. The underlined metrics indicate significantly better performance than all four baselines. **Bold** metrics indicate the overall best performance across all choices. CI is not reported for majority voting since it does not involve any randomness.

it with the standard dot product self-attention (Vaswani et al. 2017) and LRFormer (Ye et al. 2023) – a recently proposed self-attention module that prevents transformer overconfidence via Lipschitz Regularization. In addition, *UFNet* uses task-specific features and prediction scores, making it a hybrid fusion approach. We implement an early fusion approach to analyze the effect of removing task-specific predictions from *UFNet*.

The calibrated self-attention module outperformed the other two attention modules in all metrics except AUROC. While the dot product self-attention yielded 0.1% higher (non-significant) AUROC, *UFNet* was significantly better in most other metrics. LRFormer performed closely in most metrics. Finally, including task-specific predictions as additional input (hybrid fusion) enhanced performance over the early fusion approach. Since our task-specific models were trained with more subjects, they likely better understand task-specific predictions, and incorporating them improved final performance.

Withholding Uncertain Predictions

The trivial usage of MC dropout improved prediction accuracy. With dropped uncertain predictions, the best *UFNet* (hybrid fusion) model achieved $88.0 \pm 0.3\%$ accuracy, $87.1 \pm 0.3\%$ balanced accuracy, $93.0 \pm 0.2\%$ AUROC, and $81.8 \pm 0.5\%$ F_1 score – a clear improvement over our base model. It also achieved better precision, sensitivity, and specificity. However, the model could now predict $97.8 \pm 0.3\%$ of the data where it is confident enough. The model’s expected calibration error (ECE) was $5.4 \pm 0.5\%$, and the Brier score was 0.097 ± 0.002 , indicating that the model-predicted probabilities aligned well with actual disease probabilities.

Applying conformal prediction further reduced misprediction at the cost of a more frequent prediction withholding. Simultaneously applying label smoothing, Platt scaling, and conformal prediction resulted in a notable improvement in all metrics, resulting in $94.2 \pm 0.0.5\%$ accuracy, $89.9 \pm 0.7\%$ F_1 score, and $95.5 \pm 0.7\%$ AUROC. But this model could predict with enough certainty in only

$52.0 \pm 2.6\%$ cases, throwing away approximately half of the test data. Please refer to the extended version of our paper for detailed results.

Performance Across Demographic Subgroups

No significant bias in model performance was observed (see Figure 2) based on sex or race in the test set (162 individuals). The average error (i.e., misclassification) rate across female participants ($n = 85$) was $14.1 \pm 7.4\%$, compared to $6.5 \pm 5.5\%$ for male participants ($n = 77$). This difference is notable but not statistically significant (p-value = 0.11 using two-sample Z -test for proportions). The error rate was $7.63 \pm 4.79\%$ for white participants ($n = 118$) and $5.56 \pm 11.39\%$ for non-white participants ($n = 18$). This difference was also non-significant based on Fisher’s exact test (Fisher’s odd ratio = 0.71, p-value = 1.0). However, the error rate varied notably based on age subgroups. The model performed well for individuals aged 50 to 80. At the same time, the error rate was higher for those aged 30 to 50 or over 80, likely due to the under-representation of these age groups, as the dataset primarily consists of individuals who are 50-80 years old (77.1% of the entire dataset).

Validation with External Datasets

Since we introduced the first multi-task dataset for PD detection in this study, we could not compare our model to similar external datasets. However, we experimented with our model on a recently released public dataset named YouTubePD (Zhou et al. 2024). This dataset consists of 283 short YouTube videos collected from 200+ individuals (16 confirmed PD patients). Even without adapting the feature extraction pipeline for this new task, our model (*UFNet* with Smile + Speech modalities) performed comparably to state-of-the-art – 18% improvement in accuracy, 3% reduction in AUROC and F_1 score (see Table 5). This result is promising as our video features are designed for smile and pangram utterance tasks, vastly different from the videos in the YouTubePD dataset. Adapting the features for this new task should further improve performance.




Figure 2: Misclassification rate of the best *UFNet* model across demographic subgroups. The analysis is done on the test set participants. Error bars demonstrate 95% CIs.

Model	Accuracy	F ₁ score	AUROC
YouTubePD (Zhou et al. 2024)	70.61%	61.0%	87.0%
UFNet (Smile + Speech)	88.76%	58.33%	83.83%

Table 5: Summary of performance on the YouTubePD dataset. The **bold** metrics denote best performance.

Discussions and Limitations

Video analysis provides an accessible, cost-effective, and convenient means of screening for PD, particularly benefiting individuals in remote areas or low-income countries with limited access to neurological care. We carefully selected three proposed tasks for remote completion, considering feasibility and safety. While gait analysis is standard for evaluating PD (Li and Li 2022; Liu et al. 2022), it presents logistical challenges and potential risks of falling for PD patients. Sustained phonation (holding a vowel sound for as long as possible) is another speech assessment option (Tsanas et al. 2009; Vaiciukynas et al. 2017), but its analysis becomes complicated due to inconsistencies across various recording devices. All selected tasks demonstrated reasonable predictive performance in differentiating individuals with and without PD, validating our task selection.

Despite our model’s promise of accurate PD detection, harm from misprediction remains a concern. Falsely classifying someone as a PD patient may cause undue stress and economic burden, while missing individuals with PD provide a false sense of security and create delays in getting clinical care. With 78.4% recall, our model may miss some individuals with PD, so we recommend clinical follow-up if symptoms persist or worsen, even if our model does not detect PD. We also advise individuals not to consider our model’s positive prediction a definitive diagnosis. Finally, we reduce the harm of mispredictions by withholding predictions when the model is uncertain.

Our model performs consistently across sex and ethnic

subgroups, but accuracy drops for younger (30-49 years) and older (over 80 years) age groups. We find that participants from these groups are under-represented, as over 75% of our dataset participants fall within the 50-80 age range. Until a more balanced dataset is available, we recommend applying the tool only to individuals between 50 and 80. Also, due to requiring an English pangram utterance, our best model cannot be used for non-English speakers. Future work should prioritize recruiting more participants and investigating speech tasks in other languages for the worldwide deployment of our methods.

The decision threshold used in the study could be customized based on individual preferences, as it directly affects the model’s sensitivity and specificity. We used a common 0.5 threshold, but individual preferences for risk-benefit trade-offs might necessitate adjustments. For instance, some users might seek clinical evaluation even at lower probabilities (i.e., prefer high sensitivity), while others might wait for a higher likelihood before incurring healthcare costs (i.e., prefer high specificity).

Furthermore, as videos are primarily gathered unsupervised, issues such as noncompliance with task instructions and various forms of noise are common. For example, during the finger-tapping task, many participants performed fewer than the required ten taps, often with their task-performing hand obscured from view. Background noise may distort speech features, and the presence of multiple individuals in the frame could compromise smile feature extraction. Integrating post hoc quality assessment algorithms into the data collection framework could enhance data quality by identifying quality issues and prompting users to re-record tasks if needed.

A comprehensive, large-scale dataset for Parkinson’s disease is currently unavailable due to the lack of resources and clinical collaboration. This paper addresses this gap by introducing the largest multi-task video dataset in the literature and publicly releasing the anonymized features. Unfortunately, we could not gather multiple datasets of videos or extract similar features to compare model performance across different datasets since patient videos (protected health information) are not publicly accessible. Future research on privacy-preserving ways of sharing video datasets can help bridge this gap, facilitating broader comparisons and furthering research in health AI.

Conclusion

This study demonstrates the promising efficacy of machine learning models in distinguishing individuals with PD from those without PD, requiring only a computer equipped with a webcam, microphone, and internet connection. Given the shared characteristics and nuanced distinctions among movement disorders such as Parkinson’s disease, Huntington’s disease, ataxia, and Progressive Supranuclear Palsy, these findings hold significant implications. We hope the promising initial results from this research will pave the way for extending tele-neurology applications to encompass a broader range of movement disorders.

Acknowledgments

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. The first author is supported by a Google PhD fellowship.

References

- Adnan, T.; Abdelkader, A.; Liu, Z.; Hossain, E.; Park, S.; Islam, M.; and Hoque, E. 2024. A Novel Fusion Architecture for PD Detection Using Semi-Supervised Speech Embeddings. *arXiv preprint arXiv:2405.17206*.
- Adnan, T.; Islam, M. S.; Rahman, W.; Lee, S.; Tithi, S. D.; Noshin, K.; Sarker, I.; Rahman, M. S.; and Hoque, E. 2023. Unmasking Parkinson's Disease with Smile: An AI-enabled Screening Framework. *arXiv preprint arXiv:2308.02588*.
- Ali, M. R.; Hernandez, J.; Dorsey, E. R.; Hoque, E.; and McDuff, D. 2020. Spatio-temporal attention and magnification for classification of Parkinson's disease from videos collected via the internet. In *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)*, 207–214. IEEE.
- Ali, M. R.; Sen, T.; Li, Q.; Langevin, R.; Myers, T.; Dorsey, E. R.; Sharma, S.; and Hoque, E. 2021. Analyzing head pose in remotely collected videos of people with parkinson's disease. *ACM Transactions on Computing for Healthcare*, 2(4): 1–13.
- Arnab, A.; Deghani, M.; Heigold, G.; Sun, C.; Lučić, M.; and Schmid, C. 2021. Vivit: A video vision transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6836–6846.
- Atighehchian, P.; Branchaud-Charron, F.; Freyberg, J.; Pardin, R.; Schell, L.; and Pearse, G. 2022. Baal, a bayesian active learning library. <https://github.com/baal-org/baal/>. Accessed: 2024-12-27.
- Baltrušaitis, T.; Robinson, P.; and Morency, L.-P. 2016. Openface: an open source facial behavior analysis toolkit. In *2016 IEEE winter conference on applications of computer vision (WACV)*, 1–10. IEEE.
- Bandini, A.; Orlandi, S.; Escalante, H. J.; Giovannelli, F.; Cincotta, M.; Reyes-Garcia, C. A.; Vanni, P.; Zaccara, G.; and Manfredi, C. 2017. Analysis of facial expressions in parkinson's disease through video-based automatic methods. *Journal of neuroscience methods*, 281: 7–20.
- Bertasius, G.; Wang, H.; and Torresani, L. 2021. Is space-time attention all you need for video understanding? In *ICML*, volume 2, 4.
- Bologna, M.; Espay, A. J.; Fasano, A.; Paparella, G.; Hallett, M.; and Berardelli, A. 2023. Redefining bradykinesia. *Movement disorders: official journal of the Movement Disorder Society*, 38(4): 551.
- Chen, S.; Wang, C.; Chen, Z.; Wu, Y.; Liu, S.; Chen, Z.; Li, J.; Kanda, N.; Yoshioka, T.; Xiao, X.; et al. 2022. Wavlm: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing*, 16(6): 1505–1518.
- Dorsey, E.; Sherer, T.; Okun, M. S.; and Bloem, B. R. 2018. The emerging evidence of the Parkinson pandemic. *Journal of Parkinson's disease*, 8(s1): S3–S8.
- Goetz, C. G.; Tilley, B. C.; Shaftman, S. R.; Stebbins, G. T.; Fahn, S.; Martinez-Martin, P.; Poewe, W.; Sampaio, C.; Stern, M. B.; Dodel, R.; et al. 2008. 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(15): 2129–2170.
- Grishchenko, I.; Bazarevsky, V.; Zhanfir, A.; Bazavan, E. G.; Zhanfir, M.; Yee, R.; Raveendran, K.; Zhdanovich, M.; Grundmann, M.; and Sminchisescu, C. 2022. BlazePose GHUM holistic: Real-time 3D human landmarks and pose estimation. *arXiv preprint arXiv:2206.11678*.
- Gunnery, S. D.; Habermann, B.; Saint-Hilaire, M.; Thomas, C. A.; and Tickle-Degnen, L. 2016. The relationship between the experience of hypomimia and social wellbeing in people with Parkinson's disease and their care partners. *Journal of Parkinson's disease*, 6(3): 625–630.
- Ho, A. K.; Iansek, R.; Marigliani, C.; Bradshaw, J. L.; and Gates, S. 1999. Speech impairment in a large sample of patients with Parkinson's disease. *Behavioural neurology*, 11(3): 131–137.
- Hughes, A. J.; Daniel, S. E.; Kilford, L.; and Lees, A. J. 1992. Accuracy of clinical diagnosis of idiopathic Parkinson's disease: a clinico-pathological study of 100 cases. *Journal of neurology, neurosurgery & psychiatry*, 55(3): 181–184.
- Islam, M. S.; Rahman, W.; Abdelkader, A.; Lee, S.; Yang, P. T.; Purks, J. L.; Adams, J. L.; Schneider, R. B.; Dorsey, E. R.; and Hoque, E. 2023. Using AI to measure Parkinson's disease severity at home. *npj Digital Medicine*, 6(1): 156.
- Jin, B.; Qu, Y.; Zhang, L.; and Gao, Z. 2020. Diagnosing Parkinson disease through facial expression recognition: video analysis. *Journal of medical Internet research*, 22(7): e18697.
- Kissani, N.; Liqali, L.; Hakimi, K.; Mugumbate, J.; Daniel, G. M.; Ibrahim, E. A. A.; Yewnetu, E.; Belo, M.; Wilmschurst, J.; Mbelesso, P.; et al. 2022. Why does Africa have the lowest number of Neurologists and how to cover the Gap? *Journal of the Neurological Sciences*, 434: 120119.
- Langevin, R.; Ali, M. R.; Sen, T.; Snyder, C.; Myers, T.; Dorsey, E. R.; and Hoque, M. E. 2019. The PARK framework for automated analysis of Parkinson's disease characteristics. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 3(2): 1–22.
- Li, A.; and Li, C. 2022. Detecting parkinson's disease through gait measures using machine learning. *Diagnostics*, 12(10): 2404.
- Li, K.; Wang, Y.; Zhang, J.; Gao, P.; Song, G.; Liu, Y.; Li, H.; and Qiao, Y. 2023. Uniformer: Unifying convolution and self-attention for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(10): 12581–12600.

- Liu, Y.; Zhang, G.; Tarolli, C. G.; Hristov, R.; Jensen-Roberts, S.; Waddell, E. M.; Myers, T. L.; Pawlik, M. E.; Soto, J. M.; Wilson, R. M.; et al. 2022. Monitoring gait at home with radio waves in Parkinson's disease: A marker of severity, progression, and medication response. *Science Translational Medicine*, 14(663): eadc9669.
- Lugaresi, C.; Tang, J.; Nash, H.; McClanahan, C.; Uboweja, E.; Hays, M.; Zhang, F.; Chang, C.-L.; Yong, M.; Lee, J.; et al. 2019. Mediapipe: A framework for perceiving and processing reality. In *Third workshop on computer vision for AR/VR at IEEE computer vision and pattern recognition (CVPR)*, volume 2019.
- Moreau, C.; Rouaud, T.; Grabli, D.; Benatru, I.; Remy, P.; Marques, A.-R.; Drapier, S.; Mariani, L.-L.; Roze, E.; Devos, D.; et al. 2023. Overview on wearable sensors for the management of Parkinson's disease. *npj Parkinson's Disease*, 9(1): 153.
- Müller, R.; Kornblith, S.; and Hinton, G. E. 2019. When does label smoothing help? *Advances in neural information processing systems*, 32.
- Nixon, J.; Dusenberry, M. W.; Zhang, L.; Jerfel, G.; and Tran, D. 2019. Measuring calibration in deep learning. In *CVPR workshops*, volume 2.
- Platt, J.; et al. 1999. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers*, 10(3): 61–74.
- Rahman, W.; Lee, S.; Islam, M. S.; Antony, V. N.; Ratnu, H.; Ali, M. R.; Mamun, A. A.; Wagner, E.; Jensen-Roberts, S.; Waddell, E.; et al. 2021. Detecting parkinson disease using a web-based speech task: Observational study. *Journal of medical Internet research*, 23(10): e26305.
- Rufibach, K. 2010. Use of Brier score to assess binary predictions. *Journal of clinical epidemiology*, 63(8): 938–939.
- Saunders, C.; Gammernan, A.; and Vovk, V. 1999. Transduction with confidence and credibility.
- Sibley, K. G.; Girges, C.; Hoque, E.; and Foltynie, T. 2021. Video-based analyses of Parkinson's disease severity: A brief review. *J. Parkinsons. Dis.*, 11(s1): S83–S93.
- Siderowf, A.; Concha-Marambio, L.; Lafontant, D.-E.; Farris, C. M.; Ma, Y.; Urenia, P. A.; Nguyen, H.; Alcalay, R. N.; Chahine, L. M.; Foroud, T.; et al. 2023. Assessment of heterogeneity among participants in the Parkinson's Progression Markers Initiative cohort using α -synuclein seed amplification: a cross-sectional study. *The Lancet Neurology*, 22(5): 407–417.
- Tong, Z.; Song, Y.; Wang, J.; and Wang, L. 2022. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *Advances in neural information processing systems*, 35: 10078–10093.
- Tsanas, A.; Little, M.; McSharry, P.; and Ramig, L. 2009. Accurate telemonitoring of Parkinson's disease progression by non-invasive speech tests. *Nature Precedings*, 1–1.
- Vaiciukynas, E.; Verikas, A.; Gelzinis, A.; and Bacauskiene, M. 2017. Detecting Parkinson's disease from sustained phonation and speech signals. *PloS one*, 12(10): e0185613.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Yang, Y.; Yuan, Y.; Zhang, G.; Wang, H.; Chen, Y.-C.; Liu, Y.; Tarolli, C. G.; Crepeau, D.; Bukartyk, J.; Junna, M. R.; et al. 2022. Artificial intelligence-enabled detection and assessment of Parkinson's disease using nocturnal breathing signals. *Nature medicine*, 28(10): 2207–2215.
- Ye, W.; Ma, Y.; Cao, X.; and Tang, K. 2023. Mitigating transformer overconfidence via Lipschitz regularization. In *Uncertainty in Artificial Intelligence*, 2422–2432. PMLR.
- Zhou, A.; Li, S.; Sriram, P.; Li, X.; Dong, J.; Sharma, A.; Zhong, Y.; Luo, S.; Kindratenko, V.; Heintz, G.; et al. 2024. YouTubePD: A Multimodal Benchmark for Parkinson's Disease Analysis. *Advances in Neural Information Processing Systems*, 36.