

PhoneMD: Learning to Diagnose Parkinson's Disease from Smartphone Data

Patrick Schwab

Institute of Robotics and Intelligent Systems
ETH Zurich, Switzerland
patrick.schwab@hest.ethz.ch

Walter Karlen

Institute of Robotics and Intelligent Systems
ETH Zurich, Switzerland
walter.karlen@ieee.org

Abstract

Parkinson's disease is a neurodegenerative disease that can affect a person's movement, speech, dexterity, and cognition. Clinicians primarily diagnose Parkinson's disease by performing a clinical assessment of symptoms. However, misdiagnoses are common. One factor that contributes to misdiagnoses is that the symptoms of Parkinson's disease may not be prominent at the time the clinical assessment is performed. Here, we present a machine-learning approach towards distinguishing between people with and without Parkinson's disease using long-term data from smartphone-based walking, voice, tapping and memory tests. We demonstrate that our attentive deep-learning models achieve significant improvements in predictive performance over strong baselines (area under the receiver operating characteristic curve = 0.85) in data from a cohort of 1853 participants. We also show that our models identify meaningful features in the input data. Our results confirm that smartphone data collected over extended periods of time could in the future potentially be used as a digital biomarker for the diagnosis of Parkinson's disease.

Introduction

Parkinson's disease (PD) affects more than 6 million people worldwide (Vos et al. 2016) and is the second most common neurodegenerative disease after Alzheimer's disease (De Lau and Breteler 2006). The symptoms of PD progressively worsen over time, leading to a stark loss in quality of life (Schrag, Jahanshahi, and Quinn 2000), and a significant reduction in life expectancy (De Lau and Breteler 2006). While there currently exists no cure for PD, available pharmacological and surgical treatment options are effective at managing the symptoms of PD (Goetz et al. 2005; Connolly and Lang 2014). Receiving a timely and accurate diagnosis is paramount for patients because access to treatments could improve their quality of life (Global Parkinson's Disease Survey Steering Committee 2002). Currently, clinicians diagnose PD primarily based on subjective clinical assessments of patients' symptoms (Pahwa and Lyons 2010). However, research has shown that around 25% of PD diagnoses are incorrect when compared to results of post-mortem autopsy (Pahwa and Lyons 2010). Diagnosing PD is difficult because there are other movement disorders that

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Smartphones can be used to perform tests that are designed to trigger symptoms of Parkinson's disease (top). During these tests, smartphone sensors record high-resolution signals (bottom) that we can use to distinguish between people with and without Parkinson's disease.

may appear similar to PD, and because symptom severity in PD may fluctuate over time (Pahwa and Lyons 2010).

Smartphone-based tests could potentially give clinicians access to long-term measurements of symptom severity and symptom fluctuation by enabling patients to record themselves outside the clinic (Figure 1). However, making sense of observational smartphone data is extremely challenging for both humans and machines due to the large number of diverse data streams sampled at high resolution over long periods of time. Major unsolved questions include how to simultaneously cover the wide range of symptoms associated with PD, how to best aggregate the vast amounts of clinically relevant data collected over time, and how to communicate the decisions of predictive models to clinicians.

To address these issues, we present a novel approach towards distinguishing between people with and without PD from smartphone data. Our method is built on the idea of first training specialised models to assess symptom severity from single test instances, and then using an evidence aggregation model to aggregate an arbitrary number of assessments from several types of tests into a final prediction. We extend our method with hierarchical attention to visualise both the importance of tests as well as the importance of segments in those tests towards a prediction. Our experiments demonstrate that this approach leads to significant improvements in predictive performance over several strong baselines, and highlight the potential of smartphones to become accessible tools for gathering clinically relevant data in the wild.

Contributions. Our contributions are as follows:

- We present machine-learning models to assess symptoms of PD from signals recorded during smartphone-based walking, voice, tapping and memory tests.
- We introduce an evidence aggregation model (EAM) to integrate arbitrary numbers of symptom assessments from multiple types of tests over long periods of time to produce a single diagnostic score.
- We develop a hierarchical neural attention mechanism that quantifies the importance of both individual tests and segments within those tests towards the diagnostic score.
- We perform experiments on real-world data collected from 1853 mPower participants with and without PD that show that our approach leads to significant improvements in prediction performance over several strong baselines.

Related Work

Background. Machine learning has a rich history in facilitating medical diagnoses. Machine learning has, for example, been applied to diagnosing breast cancer from tumor features (Zheng, Yoon, and Lam 2014), cardiac arrhythmias and cardiac risk factors from smartphone-based heart rate sensors (Oresko et al. 2010; Schwab et al. 2017; Ballinger et al. 2018), skin cancer from clinical images (Esteve et al. 2017), depressed moods from information self-reported via smartphones (Suhara, Xu, and Pentland 2017), and a wide range of clinical diagnosis codes from electronic health records and lab test results (Lipton et al. 2015; Choi et al. 2016; Razavian, Marcus, and Sontag 2016). Predicting a person's disease status is difficult because there is a vast range of factors that may influence an individual's health. Wearable sensors and smart devices enable us to capture a number of these factors with minimal burden on users by passively and continuously tracking behaviors and environmental factors (Quisel et al. 2017). However, in contrast to clean, standardised benchmark datasets, observational data collected by wearable sensors and smart devices in the real-world is often difficult to integrate with existing machine-learning approaches. The difficulty of applying existing machine-learning methods to complex datasets has led to the development of specialised methods to deal with several of the idiosyncrasies of observational health data, such as missingness (Lipton, Kale, and Wetzel 2016; Che et al. 2018), long-term temporal dependencies (Choi et al. 2016), noise (Schwab et al. 2017), heterogeneity (Libbrecht and Noble 2015), irregular sampling (Lipton et al. 2015), sparsity (Lasko, Denny, and Levy 2013), and multivariate input data (Ghassemi et al. 2015; Schwab et al. 2018). However, adapting existing machine-learning methods to account for the idiosyncrasies of healthcare data remains an ongoing challenge (Ghassemi et al. 2018).

Monitoring and Diagnosis of PD. There has been much interest in leveraging new technologies and data modalities to better diagnose and assess symptom severity in PD. There are a number of driving factors behind the interest in new approaches: Firstly, despite the severity of the disease, clinical PD diagnoses are currently relatively inaccurate. Diagnoses

are particularly difficult in the earlier stages of the disease and in the presence of other disorders that may appear similar to PD (Rizzo et al. 2016). Secondly, new technologies could lead to patients receiving their diagnoses earlier. An early diagnosis could potentially improve a patient's quality of life by giving them access to symptom-suppressing treatments (Global Parkinson's Disease Survey Steering Committee 2002). Lastly, both clinical trials for new pharmaceutical treatments and clinical decision-making require the ability to accurately diagnose and objectively assess symptoms of PD (Shulman et al. 2006; Dorsey et al. 2017). Previous works have for example used data from pen movements (Smith et al. 2007), wearables (Patel et al. 2009; Klucken et al. 2013), and speech features (Little et al. 2007; 2009; Tsanas et al. 2010; 2011; 2012) to objectively monitor or diagnose PD. A number of works have also proposed the use of smartphone sensors for continuously monitoring symptoms in PD (Hammerla et al. 2015; Arora et al. 2015; Zhan et al. 2016; 2018; Prince, Andreotti, and Vos 2018). Recently, the PD Digital Biomarker DREAM challenge¹ aimed to develop machine-learning models to diagnose PD from accelerometer data in a collaborative effort. (Emrani, McGuirk, and Xiao 2017) proposed a multitask-learning framework to identify biomarkers that are predictive of progression in PD. However, their approach did not integrate raw sensor data and could not handle missing input data.

In contrast to existing works, we present the first machine-learning approach to distinguishing between people with and without PD that integrates information from sensor measurements of several types of smartphone-based tests over long periods of time. Our approach is able to simultaneously (i) assess single test instances and (ii) produce a unified diagnostic score. In addition, we introduce a hierarchical neural attention mechanism that enables us to reason about both the importance of specific tests as well as the importance of individual segments within those tests towards the final diagnostic score. Furthermore, we perform our experiments on data collected from 1853 mPower participants, the largest cohort used to validate a machine-learning approach to diagnosing PD from smartphone data to date.

Methodology

Overview. We utilise data collected during the mPower study, a large-scale observational study about PD conducted entirely through a smartphone app (Bot et al. 2016). In the study, participants with and without PD are asked to perform four smartphone-based tests (walking, voice, tapping and memory; Figure 1) up to three times a day without any supervision. In addition to regularly performing the tests, participants provide their detailed demographic profile, including possible prior clinical diagnoses of PD, using self-reporting forms within the app². The main idea of the presented approach is to connect the sensor data collected by the participants' smartphones with their prior professional diagnoses to train machine-learning models to learn to diagnose PD.

¹<http://synapse.org/DigitalBiomarkerChallenge>

²<https://parkinsonmpower.org/>




Figure 2: An illustration of the data processing pipelines for each of the test types (vertical) from the input signals x_* (left) over the specialised predictive models P_* (center) to the single-test output predictions y_* (right). The use of specialised predictive models for each test type enables us to choose the most suitable model for each of the heterogeneous input signals.

Smartphone Tests. mPower participants perform the following four types of tests using their personal smartphones:

Walking Test. To perform the walking test, participants are asked to put their smartphone in their pocket, walk 20 steps forward, turn around, stand still for 30 seconds, and then walk 20 steps back. We denote the three distinct segments of the walking test as: Outbound, rest, and return, respectively. During the test, the smartphone’s accelerometer and gyroscope record the participant’s three-dimensional linear and angular acceleration. This test is designed to measure movement impairments associated with PD, such as tremor, rigidity, and freezing of gait.

Voice Test. In the voice test, participants are asked to say “aaaah” into their smartphones’ microphone for up to 10 seconds. The smartphone’s microphone records the audio data during the test and during the preceding countdown. The goal of the audio test is to expose speech impairments that are commonly found in people with PD.

Tapping Test. In the tapping test, participants are asked to position their smartphones on a flat surface and alternately tap two buttons on the screen for 20 seconds. The smartphone records the positions and timestamps of the participant’s taps on the screen. In addition, the smartphone’s accelerometer measures the three-dimensional movements of the smartphone during the test. The tapping test is aimed at uncovering signs of impaired finger dexterity. Impaired finger dexterity is a common symptom in people with PD.

Memory Test. In the memory test, participants are presented with a grid of flowers on their smartphone screens. During the test, different flowers are illuminated one at a time. Participants are then asked to repeat the observed sequence by touching the flowers in the same order. The collected data includes the positions and timestamps of the participant’s taps on the smartphone’s screen and the sequence order as displayed to the participant. This test measures the spatial memory of the participant, which may be impaired due to PD (Bot et al. 2016).




Figure 3: Temporal ensembling using an evidence aggregation model (EAM). The EAM (grey) receives per-test meta-data $(m_{*,j})$ and per-test output predictions $(y_{*,j})$ in temporal order as input. In this example, the EAM’s hidden state (h_j) aggregates the information from the $k = 4$ performed tests to produce a final output y that indicates whether or not the participant is likely to have PD.




Multistage Approach. Our approach to distinguishing between people with and without PD consists of two stages. In the first stage, we use specialised predictive models P_* to identify PD in signals x_* from a single type of test with $\star \in \{\text{Walking}, \text{Voice}, \text{Tapping}, \text{Memory}\}$. The specialised models are trained to predict a participant’s diagnosis given the signal data from exactly one sample instance of one type of test (Figure 2). The output of the specialised models is a local prediction y_* that indicates, on a scale from 0 to 1, how likely it is that the participant that performed the given test instance has PD:

$$y_* = P_*(x_*) \quad (1)$$

The specialised models P_* are the building blocks for the second stage. In the second stage, the outputs $y_{*,j}$, with $j \in [1, k]$, of the specialised models and the metadata $m_{*,j}$ for all k tests performed by a user are aggregated into a single diagnostic prediction y using an EAM (Figure 3):

$$y = \text{EAM}([(m_{*,1}, y_{*,1}), \dots, (m_{*,k}, y_{*,k})]) \quad (2)$$

The primary idea behind Equations 1 and 2 is to disentangle learning how to assess symptom severity from each test and how to aggregate multiple tests over a period of time. This compositional approach to modelling the problem of diagnosing PD from a range of diverse smartphone tests enables us to choose the most suitable predictive model for the various test types and the EAM. Furthermore, each specialised predictive model P_* is optimised for one type of test only. In contrast to an end-to-end model, the specialised predictive models do not need to consider how to aggregate multiple tests and which patterns may be important in other tests. Similarly, the EAM is entirely focused on learning how to best aggregate the evidence from multiple tests. In essence, our approach follows a divide-and-conquer approach by ensuring that each component is focused on exactly one task. Another benefit of the given abstract formulation is that it enables us to choose from a wide range of models for both the specialised predictive models and the EAM, since there are no specific requirements on either other than that they need to process x_* and tuples of $(m_{*,i}, y_{*,i})$, respectively.

	 outbound		 rest		 return	
	CNN	Feature	CNN	Feature	CNN	Feature
AUC	0.53 (0.50, 0.56)	0.50 (0.50, 0.53)	0.53 (0.50, 0.56)	0.52 (0.50, 0.55)	0.77 (0.74, 0.79)	0.77 (0.75, 0.79)
AUPR	0.60 (0.57, 0.64)	0.60 (0.55, 0.62)	0.62 (0.59, 0.66)	0.61 (0.55, 0.62)	0.72 (0.53, 0.87)	0.86 (0.84, 0.88)




	 voice		 tapping		 memory	
	CNN	Feature	CNN	Feature	RNN	Feature
AUC	0.53 (0.50, 0.55)	0.56 (0.54, 0.58)	0.59 (0.57, 0.61)	0.62 (0.60, 0.64)	0.65 (0.60, 0.69)	0.52 (0.50, 0.57)
AUPR	0.48 (0.45, 0.51)	0.45 (0.43, 0.48)	0.56 (0.53, 0.59)	0.65 (0.62, 0.67)	0.91 (0.88, 0.93)	0.87 (0.84, 0.89)

Table 1: Comparison of the AUC and AUPR values for the different test types when only given the data of a single test to make a diagnostic decision. We compared the performances of neural networks (CNN, RNN) with expert features from biomedical literature fed to a random forest model (Feature) on the validation set. The listed models were the best models encountered over 35 hyperparameter optimisation runs for each test and model type. We calculated the 95% confidence intervals (CIs) using bootstrap resampling with 1000 bootstrap samples. A comparison between the test types was not possible, because the evaluated subsets differed significantly due to different user groups preferring to do certain tests in different amounts (Appendix D).

Hierarchical Neural Attention. In addition to the diagnostic score y , our approach provides the clinician with information about which tests and test segments in the data recorded by the user were most important for the model’s output. Presenting information about which data the model output is based on can help put the diagnostic score y in perspective and inform the clinician’s further clinical decision-making. For example, in a patient whose diagnostic prediction focused primarily on motor symptoms, the clinician can focus her efforts on ruling out other movement disorders that may cause similar symptoms. In order to highlight (i) which individual tests were most important for the EAM’s output y , and (ii) which segments of specific tests were most important for the local predictions y_* , we introduce a hierarchical neural soft attention mechanism. When using neural networks as predictive models, the upper-level attention mechanism (i) is a component of the EAM and the lower-level attention mechanism (ii) is part of the specialised models P_* . Both the upper- and lower-level attention mechanism use the same mathematical formulation. Given the top-most hidden feature representations h_i of (i) all the tests performed by a user, or (ii) segments in the recorded signal streams for a single test, we calculate attention factors a_i using (Xu et al. 2015; Schwab et al. 2017; Schwab, Miladinovic, and Karlen 2019):

$$a_i = \frac{\exp(u_i^T u_s)}{\sum_{j=1}^m \exp(u_j^T u_s)} \quad (3)$$

where

$$u_i = \text{activation}(W_s h_i + b_s) \quad (4)$$

Equation (4) corresponds to a single-layer MLP with a weight matrix W_s and bias b_s . The single-layer MLP projects h_i into a suitable hidden representation u_i for comparison with u_s . We then calculate the attention factors a_i by computing the softmax similarity of u_i to u_s . u_s is the most informative hidden representation, i.e. the hidden representation for which a_i would be the highest (Schwab, Miladinovic, and Karlen 2019). W_s , b_s and u_s are learned parameters and jointly optimised with the other parameters during training.

Experiments

Our experiments aimed to answer the following questions:

- 1 What is the comparative performance of various specialised models P_* in diagnosing PD from a single test?
- 2 How do EAMs compare to existing methods for aggregating multiple local predictions?
- 3 What is the overall diagnostic accuracy of our approach?
- 4 Does the proposed hierarchical neural attention mechanism identify meaningful data points?

To answer these questions, we performed experimental comparisons between various baselines, predictive models and EAMs both on predicting PD from a single test and from an arbitrary number of tests.

Dataset and Study Cohort. We use data from the mPower study, a worldwide observational study about PD conducted entirely through smartphones (Bot et al. 2016). Starting in March 2015, the study recruited participants aged 18 and older around the world through a mobile app. Participants provided their demographic profile, including prior diagnoses of PD, through self-reporting, and performed the four test types regularly. Out of the study cohort, we used the subset of participants that were 45 or older, because there were very few participants in the dataset that had a clinical diagnosis at younger age. We used only those tests that were performed off medication, except for the memory tests. We performed a random split stratified by participant age to divide the available dataset into a training set (70%), validation set (10%), and test set (20%). Each participant and the tests they performed were assigned to exactly one of the three folds without any overlap (Table).

Models. For each test type, we trained several specialised predictive models P_* using both automated feature extraction with neural networks and random forest (RF) models. We used expert features from biomedical literature that have been shown to be predictive of PD in the given data modalities as inputs to the RF models. The complete list of features used for each test type can be found in Appendix A. For the

Method	AUC	AUPR	F1	Sens@95%Spec
EAM (Both) + age + gender	0.85 (0.81, 0.89)	0.87 (0.82, 0.91)	0.81 (0.75, 0.85)	0.43 (0.19, 0.54)
EAM (Neural networks) + age + gender	0.84 (0.80, 0.88)	0.86 (0.81, 0.90)	0.82 (0.74, 0.86)	0.33 (0.21, 0.51)
EAM (Feature) + age + gender	0.84 (0.79, 0.88)	0.86 (0.81, 0.90)	0.76 (0.73, 0.84)	0.40 (0.23, 0.56)
End-to-end neural network + age + gender	0.50 (0.50, 0.56)	0.54 (0.46, 0.62)	0.27 (0.20, 0.70)	0.04 (0.01, 0.07)
Age + gender	0.74 (0.69, 0.79)	0.75 (0.68, 0.82)	0.72 (0.67, 0.79)	0.16 (0.09, 0.31)
EAM (Both)	0.70 (0.64, 0.75)	0.74 (0.66, 0.79)	0.67 (0.60, 0.71)	0.23 (0.15, 0.41)
EAM (Neural networks)	0.71 (0.65, 0.76)	0.75 (0.67, 0.80)	0.67 (0.61, 0.72)	0.24 (0.14, 0.41)
EAM (Feature)	0.71 (0.66, 0.76)	0.75 (0.67, 0.80)	0.68 (0.61, 0.73)	0.24 (0.14, 0.39)
Mean Aggregation (Neural networks)	0.64 (0.58, 0.69)	0.67 (0.58, 0.73)	0.60 (0.52, 0.68)	0.22 (0.10, 0.27)
Mean Aggregation (Feature)	0.62 (0.56, 0.68)	0.60 (0.51, 0.66)	0.62 (0.53, 0.69)	0.13 (0.00, 0.19)
Max Aggregation (Neural networks)	0.61 (0.55, 0.67)	0.61 (0.53, 0.68)	0.59 (0.54, 0.68)	0.03 (0.01, 0.19)
Max Aggregation (Feature)	0.61 (0.54, 0.66)	0.61 (0.52, 0.68)	0.60 (0.52, 0.65)	0.07 (0.03, 0.18)

Table 2: Comparison of the AUC, AUPR, F1, and sensitivity at a fixed specificity of 95% (Sens@95%Spec) on the test set of 347 participants across the methods that we evaluated. In parentheses are the 95% CIs calculated with 1000 bootstrap samples.

neural networks, we used different architectures of neural networks for each test depending on the type of input signal. For the walking, voice and tapping task, we used multi-layer convolutional neural networks (CNNs) with max pooling and temporal convolutions. For the memory test, we used a recurrent neural network (RNN) with bidirectional long short-term memory (BLSTM). Except for the voice test, the neural networks hosted the segment-level neural attention mechanisms described previously. For the voice CNN, we did not employ a neural attention mechanism because we found that it was detrimental to predictive performance. To implement the previously described EAM, we used a RNN architecture consisting of BLSTM cells. We trained EAMs using (1) only the RF models, (2) only the neural networks, and (3) an ensemble of both as specialised models to compare the performances of both approaches and whether their outputs are complementary. The detailed architectures for the neural networks and EAM are given in Appendix B. The EAM received a one-hot encoded unique identifier of the specialised predictive model as input metadata $m_{*,j}$ with each local per-test prediction y_* . The unique identifier enabled the EAM to differentiate between the various specialised predictive models. We additionally tested passing timing information, including the time since the last per-

formed test and the hour of day at which the test was performed, for each performed test. However, we found no performance benefit in adding timing information to the metadata. Lastly, in order to determine whether the use of an EAM improves performance over simpler approaches, we evaluated the performances of aggregating over local predictions y_* using the mean and maximum values of all local predictions. As a simple baseline based on demographics, we trained a MLP that received as input the age and gender of a participant and no information of any of their performed tests. To determine whether the proposed separation of learning to assess single test instances and learning to integrate multiple tests is beneficial, we also trained an end-to-end neural network jointly on both tasks. The end-to-end neural network used the same architectures as the specialised models to assess the tests and the same architecture as the EAM to integrate multiple tests.

Hyperparameters. We took a systematic approach to hyperparameter search. To avoid biases stemming from using different degrees of hyperparameter optimisation, we evaluated exactly 35 hyperparameter configurations for each trained specialised predictive model and EAM. We report the performances of those models which achieved the best validation set performances across the 35 runs. We selected the hyperparameters at random from a fixed range for each hyperparameter run. For the RF models, we used 512 to 1024 trees in the forest and a maximum tree depth between 3 and 5. For all neural networks, we used dropout of 0 to 70% between hidden layers, an L_2 penalty of 0, 0.0001 or 0.00001, and varying numbers of layers and hidden units depending on the test type (Appendix C). For the EAM, we used 2 to 5 stacked BLSTM layers with 16 to 64 hidden units each. We optimised the neural networks’ binary cross-entropy for up to 500 epochs with a learning rate of 0.0001, a batch size of 32, and an early stopping patience of 12 epochs on the validation set. For memory reasons, we used a batch size of 2 for the end-to-end trained neural network. All other hyperparameters and hyperparameter ranges were exactly the same as in the separated models.

Property	Training	Validation	Test
Subjects (#)	1314 (70%)	192 (10%)	347 (20%)
PD (%)	52.36	50.00	56.20
Female (%)	28.00	36.98	25.94
Age (years)	59.29 \pm 9.40	59.53 \pm 9.03	58.90 \pm 9.24
Walking (#)	13.89 \pm 35.07	15.58 \pm 33.90	14.03 \pm 45.20
Voice (#)	16.11 \pm 40.21	19.47 \pm 44.55	14.88 \pm 45.12
Tapping (#)	15.20 \pm 38.04	18.50 \pm 43.12	14.78 \pm 42.67
Memory (#)	14.01 \pm 33.30	20.78 \pm 35.92	17.58 \pm 38.11
Usage (days)	24.27 \pm 41.01	29.66 \pm 45.73	25.43 \pm 43.24

Table 3: Population statistics of the training, validation, and test set. Numbers (#) shown are mean \pm standard deviation.




Figure 4: The outputs of the employed hierarchical neural attention mechanism on a user with PD that performed 18 tests. The timeline (top, left to right) shows all the tests performed by the user in temporal order. The tests performed (top, a_{test}) and the data segments within the tests (center, a_{seg}) were assigned attention weights that quantify their relative importance towards the final diagnostic score y . We show the outbound accelerometer data (left) and the rest accelerometer data (right) from walking test $\mathbf{\hat{A}}_{10}$. In the outbound recording, the attention mechanism focused strongly on the long delay to start moving (segment 1), increasingly choppy movement while setting foot (segments 3, 4, and 5), and the abrupt stop (segment 7). In the rest recording, we found that the attention was evenly distributed across the recording, likely because the whole recording contained signs of what could have been resting tremor. Slightly more attention was paid to segments with increased motion (segments 5 and 13).

Preprocessing. For computational performance reasons, we downsampled the input signals for the walking, voice and tapping test by factors of 10, 4, and 10, respectively. In our initial evaluation, we did not see significant differences in predictive performance when using higher resolution data. After downsampling, we zero-padded or truncated the size of the sensor data to fixed lengths for each test type if they were too short or too long, respectively. The fixed lengths were 300, 250, 25, and 300 samples per record for the walking, voice, memory and tapping tests, respectively. For the voice test, we did not pass the raw voice signal to the neural networks. Instead, we passed the Mel-frequency cepstral coefficients (MFCC) that we extracted from the audio signal using a window size of 30 ms, a stride of 10 ms and 40 coefficients as input signal. For the RFs, we used the raw audio signals downsampled from their original sampling rate of 44100 Hz with factor 20 as inputs to the feature extractors. We standardised the accelerometer data for the walking and tapping tests to have zero mean and unit variance.

Metrics. We computed the area under the receiver operating characteristic curve (AUC) and the area under the precision recall curve (AUPR) when comparing the different specialised predictive models. We evaluated the specialised models on the validation set to avoid test set leakage that could affect the evaluation of the models that aggregate information from multiple tests. We chose the best-performing specialised predictive models for each test for use in the aggregation models based on validation set performance. To compare the models that aggregated all available tests of a single user into a single diagnostic score, we additionally calculated the F1 score and the sensitivity at a fixed specificity level of 95%. Some of the data folds were not balanced between people with and without PD. In particular, comparing single-test performances between test types was not possible due to differences in the number of tests performed between people with and without PD (Appendix D). We evaluated the performances of the three parts of the walking test (outbound, rest, and return) separately to determine their relative importances for diagnosing PD.

Results

Single-test Performance. In terms of single-test performance, we found that, generally, both RFs with expert features and automated feature extraction with neural networks achieved competitive results for all tests (Table 1). The performances of RFs with expert features and neural networks were similar across all tests, except for the tapping test, where the expert features significantly outperformed the neural networks, and the memory test, where the neural networks likewise outperformed the expert features. When comparing the three segments of the walking test, we found that return was the most informative for diagnosing PD.

Overall Performance. We found large differences in performance between the various aggregation models that took into account all the performed tests of a user (Table 2). Notably, EAMs outperformed all baselines by a large margin, and significantly improved upon the demographics-only model by integrating information from the tests performed by participants. We also found that expert features and neural network features were to some degree complementary, as the best EAM using both sets of predictive models outperformed its counterparts that only used one set of specialised predictive models. The neural networks trained end-to-end to simultaneously assess all types of tests and aggregate information from the available tests over time failed to converge. Closer analysis revealed that the end-to-end network was unable to effectively propagate gradients through the initially more attractive upper layers down to the per-test layers. Disentangling symptom assessment and temporal aggregation enabled EAMs to overcome this issue entirely.

Hierarchical Attention. We plotted the attributions of the hierarchical neural attention mechanism against the raw signals of a sample participant with PD (Figure 4). In the walking tests, the attributions potentially corresponded to regions where signs of resting tremor and rigid motions could have appeared. In the memory tests, we found that the focus was directed at the difficult end stage of the test (Appendix E).

Discussion

Our work expands on prior studies (Arora et al. 2015) by developing an effective methodology for integrating evidence from multiple types of smartphone-based tests over long periods of time, introducing tools to identify the most salient data segments across the vast amounts of generated data points, and evaluating these novel approaches in a large, representative cohort. The availability of smartphone-based tools for diagnosing PD could have a profound impact on clinical practice by enabling clinicians to access long-term observational data on patients. These additional data points could help give clinicians a more comprehensive and objective view on their patients' symptoms and symptom fluctuations, and therefore possibly enable more accurate diagnoses and treatment regimes. Another major potential benefit of enabling patients to record themselves with their smartphones is that it could enable clinicians to monitor their patients without requiring in-person visits that may be time-consuming and expensive, particularly in rural locations and developing countries. While our initial results are promising, further clinical validation is needed to determine whether the availability of smartphone data, the proposed diagnostic score, and in-depth information about the most relevant data points improve clinicians' ability to accurately diagnose PD.

Limitations. The main limitation of this work is that we use prior clinical diagnoses of users to train and evaluate our models. Clinical diagnoses for PD are themselves often inaccurate (Rizzo et al. 2016), and are therefore not a flawless gold standard to evaluate against. In addition, much like in clinical assessments, smartphone-based tests depend on PD symptoms being clearly distinguishable for at least some of the tests being performed. While smartphones enable patients to record themselves when they believe that their symptoms are most pronounced, they still might not be clearly distinguishable against normal human variability, particularly in early-onset PD. Furthermore, the accuracy of smartphone diagnostics may be reduced when confronted with other movement and neurologic disorders that may appear similar to PD. More data, ideally from a prospective study, is needed to conclusively determine the robustness of machine-learning and smartphone-based tests against these confounding factors.

Conclusion

We presented a machine-learning approach to distinguishing between people with and without PD from multiple smartphone-based tests. Our multistage approach is built on the idea of separately training (i) specialised models to assess symptom severity in instances of a single test, and (ii) an EAM to integrate all available single-test assessments into a final diagnostic score. In addition, we introduced a hierarchical attention mechanism that shows both which tests out of all performed tests, and which segments within those tests were most important for the model's decision. We demonstrated experimentally that the presented approach leads to significant improvements over several strong baselines with an AUC of 0.85 (95% CI: 0.81, 0.89), an AUPR of 0.87

(95% CI: 0.82, 0.91) and a sensitivity at 95% specificity of 43% (95% CI: 0.19, 0.54) in data from a cohort of 1853 participants. Our results confirm that machine-learning algorithms and smartphone data collected in the wild over extended periods of time could in the future potentially be used as a digital biomarker for the diagnosis of PD.

Acknowledgments. This work was partially funded by the Swiss National Science Foundation (SNSF) projects 167302 and 150640. We acknowledge the support of NVIDIA Corporation with the donation of the GPUs used for this research. The data used in this manuscript were contributed by users of the Parkinson mPower mobile application as part of the mPower study developed by Sage Bionetworks and described in Synapse (doi:10.7303/syn4993293).

References

- Arora, S.; Venkataraman, V.; Zhan, A.; Donohue, S.; Biglan, K.; Dorsey, E.; and Little, M. 2015. Detecting and monitoring the symptoms of Parkinson's disease using smartphones: A pilot study. *Parkinsonism & related disorders* 21(6):650–653.
- Ballinger, B.; Hsieh, J.; Singh, A.; Sohoni, N.; Wang, J.; Tison, G. H.; Marcus, G. M.; Sanchez, J. M.; Maguire, C.; Olgin, J. E.; et al. 2018. DeepHeart: Semi-Supervised Sequence Learning for Cardiovascular Risk Prediction. In *AAAI Conference on Artificial Intelligence*.
- Bot, B. M.; Suver, C.; Neto, E. C.; Kellen, M.; Klein, A.; Bare, C.; Doerr, M.; Pratap, A.; Wilbanks, J.; Dorsey, E. R.; et al. 2016. The mPower study, Parkinson disease mobile data collected using ResearchKit. *Scientific data* 3:160011.
- Che, Z.; Purushotham, S.; Cho, K.; Sontag, D.; and Liu, Y. 2018. Recurrent neural networks for multivariate time series with missing values. *Scientific reports* 8(1):6085.
- Choi, E.; Bahadori, M. T.; Schuetz, A.; Stewart, W. F.; and Sun, J. 2016. Doctor AI: Predicting clinical events via recurrent neural networks. In *Machine Learning for Healthcare Conference*.
- Connolly, B. S., and Lang, A. E. 2014. Pharmacological treatment of Parkinson disease: a review. *JAMA* 311(16):1670–1683.
- De Lau, L. M., and Breteler, M. M. 2006. Epidemiology of Parkinson's disease. *The Lancet Neurology* 5(6):525–535.
- Dorsey, E. R.; Papapetropoulos, S.; Xiong, M.; and Kiebertz, K. 2017. The first frontier: digital biomarkers for neurodegenerative disorders. *Digital Biomarkers* 1(1):6–13.
- Emrani, S.; McGuirk, A.; and Xiao, W. 2017. Prognosis and Diagnosis of Parkinson's Disease Using Multi-Task Learning. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1457–1466. ACM.
- Esteva, A.; Kuprel, B.; Novoa, R. A.; Ko, J.; Swetter, S. M.; Blau, H. M.; and Thrun, S. 2017. Dermatologist-level classification of skin cancer with deep neural networks. *Nature* 542(7639):115.
- Ghassemi, M.; Pimentel, M. A.; Naumann, T.; Brennan, T.; Clifton, D. A.; Szolovits, P.; and Feng, M. 2015. A Multivariate Timeseries Modeling Approach to Severity of Illness Assessment and Forecasting in ICU with Sparse, Heterogeneous Clinical Data. In *AAAI Conference on Artificial Intelligence*.
- Ghassemi, M.; Naumann, T.; Schulam, P.; Beam, A. L.; and Ranganath, R. 2018. Opportunities in Machine Learning for Healthcare. *arXiv preprint arXiv:1806.00388*.
- Global Parkinson's Disease Survey Steering Committee. 2002. Factors impacting on quality of life in Parkinson's disease: results from an international survey. *Movement Disorders* 17(1):60–67.

- Goetz, C. G.; Poewe, W.; Rascol, O.; and Sampaio, C. 2005. Evidence-based medical review update: pharmacological and surgical treatments of Parkinson's disease: 2001 to 2004. *Movement Disorders* 20(5):523–539.
- Hammerla, N. Y.; Fisher, J.; Andras, P.; Rochester, L.; Walker, R.; and Plötz, T. 2015. PD Disease State Assessment in Naturalistic Environments Using Deep Learning. In *AAAI Conference on Artificial Intelligence*.
- Klucken, J.; Barth, J.; Kugler, P.; Schlachetzki, J.; Henze, T.; Marxreiter, F.; Kohl, Z.; Steidl, R.; Hornegger, J.; Eskofier, B.; et al. 2013. Unbiased and mobile gait analysis detects motor impairment in Parkinson's disease. *PLoS one* 8(2):e56956.
- Lasko, T. A.; Denny, J. C.; and Levy, M. A. 2013. Computational phenotype discovery using unsupervised feature learning over noisy, sparse, and irregular clinical data. *PLoS one* 8(6):e66341.
- Libbrecht, M. W., and Noble, W. S. 2015. Machine learning applications in genetics and genomics. *Nature Reviews Genetics*.
- Lipton, Z. C.; Kale, D. C.; Elkan, C.; and Wetzel, R. 2015. Learning to diagnose with LSTM recurrent neural networks. In *International Conference on Learning Representations*.
- Lipton, Z. C.; Kale, D. C.; and Wetzel, R. 2016. Directly modeling missing data in sequences with RNNs: Improved classification of clinical time series. In *Machine Learning for Healthcare Conference*.
- Little, M. A.; McSharry, P. E.; Roberts, S. J.; Costello, D. A.; and Moroz, I. M. 2007. Exploiting nonlinear recurrence and fractal scaling properties for voice disorder detection. *BME Online* 6(1):23.
- Little, M. A.; McSharry, P. E.; Hunter, E. J.; Spielman, J.; Ramig, L. O.; et al. 2009. Suitability of dysphonia measurements for telemonitoring of Parkinson's disease. *IEEE Transactions on Biomedical Engineering* 56(4):1015–1022.
- Oresko, J. J.; Jin, Z.; Cheng, J.; Huang, S.; Sun, Y.; Duschl, H.; and Cheng, A. C. 2010. A wearable smartphone-based platform for real-time cardiovascular disease detection via electrocardiogram processing. *IEEE Transactions on Information Technology in Biomedicine* 14(3):734–740.
- Pahwa, R., and Lyons, K. E. 2010. Early diagnosis of Parkinson's disease: recommendations from diagnostic clinical guidelines. *Am J Manag Care* 16(4):94–99.
- Patel, S.; Lorincz, K.; Hughes, R.; Huggins, N.; Growdon, J.; Standaert, D.; Akay, M.; Dy, J.; Welsh, M.; and Bonato, P. 2009. Monitoring motor fluctuations in patients with Parkinson's disease using wearable sensors. *IEEE Transactions on Information Technology in Biomedicine* 13(6):864–873.
- Prince, J.; Andreotti, F.; and Vos, M. D. 2018. Multi-Source ensemble learning for the remote prediction of Parkinson's disease in the presence of source-wise missing data. *IEEE Transactions on Biomedical Engineering*.
- Quisel, T.; Foschini, L.; Signorini, A.; and Kale, D. C. 2017. Collecting and analyzing millions of mHealth data streams. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1971–1980. ACM.
- Razavian, N.; Marcus, J.; and Sontag, D. 2016. Multi-task prediction of disease onsets from longitudinal laboratory tests. In *Machine Learning for Healthcare Conference*.
- Rizzo, G.; Copetti, M.; Arcuti, S.; Martino, D.; Fontana, A.; and Logroscino, G. 2016. Accuracy of clinical diagnosis of Parkinson disease: A systematic review and meta-analysis. *Neurology*.
- Schrag, A.; Jahanshahi, M.; and Quinn, N. 2000. How does Parkinson's disease affect quality of life? A comparison with quality of life in the general population. *Movement Disorders*.
- Schwab, P.; Scebba, G. C.; Zhang, J.; Delai, M.; and Karlen, W. 2017. Beat by Beat: Classifying Cardiac Arrhythmias with Recurrent Neural Networks. In *Computing in Cardiology*.
- Schwab, P.; Keller, E.; Muroi, C.; Mack, D. J.; Strässle, C.; and Karlen, W. 2018. Not to Cry Wolf: Distantly Supervised Multitask Learning in Critical Care. In *International Conference on Machine Learning*.
- Schwab, P.; Miladinovic, D.; and Karlen, W. 2019. Granger-causal Attentive Mixtures of Experts: Learning Important Features with Neural Networks. In *AAAI Conference on Artificial Intelligence*.
- Shulman, L. M.; Pretzer-Abhoff, I.; Anderson, K. E.; Stevenson, R.; Vaughan, C. G.; Gruber-Baldini, A. L.; Reich, S. G.; and Weiner, W. J. 2006. Subjective report versus objective measurement of activities of daily living in Parkinson's disease. *Movement Disorders*.
- Smith, S. L.; Gaughan, P.; Halliday, D. M.; Ju, Q.; Aly, N. M.; and Playfer, J. R. 2007. Diagnosis of Parkinson's disease using evolutionary algorithms. *Genetic Programming and Evolvable Machines* 8(4):433–447.
- Suhara, Y.; Xu, Y.; and Pentland, A. 2017. Deepmood: Forecasting depressed mood based on self-reported histories via recurrent neural networks. In *International Conference on World Wide Web*.
- Tsanas, A.; Little, M. A.; McSharry, P. E.; and Ramig, L. O. 2010. Accurate telemonitoring of Parkinson's disease progression by noninvasive speech tests. *IEEE Transactions on Biomedical Engineering* 57(4):884–893.
- Tsanas, A.; Little, M. A.; McSharry, P. E.; and Ramig, L. O. 2011. Nonlinear speech analysis algorithms mapped to a standard metric achieve clinically useful quantification of average Parkinson's disease symptom severity. *Journal of the Royal Society Interface*.
- Tsanas, A.; Little, M. A.; McSharry, P. E.; Spielman, J.; and Ramig, L. O. 2012. Novel speech signal processing algorithms for high-accuracy classification of Parkinson's disease. *IEEE Transactions on Biomedical Engineering* 59(5):1264–1271.
- Vos, T.; Allen, C.; Arora, M.; Barber, R. M.; Bhutta, Z. A.; Brown, A.; Carter, A.; Casey, D. C.; Charlson, F. J.; Chen, A. Z.; et al. 2016. Global, regional, and national incidence, prevalence, and years lived with disability for 310 diseases and injuries, 1990–2015: a systematic analysis for the global burden of disease study 2015. *The Lancet* 388(10053):1545–1602.
- Xu, K.; Ba, J.; Kiros, R.; Cho, K.; Courville, A.; Salakhudinov, R.; Zemel, R.; and Bengio, Y. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *International Conference on Machine Learning*, 2048–2057.
- Zhan, A.; Little, M. A.; Harris, D. A.; Abiola, S. O.; Dorsey, E.; Saria, S.; and Terzis, A. 2016. High Frequency Remote Monitoring of Parkinson's Disease via Smartphone: Platform Overview and Medication Response Detection. *arXiv preprint arXiv:1601.00960*.
- Zhan, A.; Mohan, S.; Tarolli, C.; Schneider, R. B.; Adams, J. L.; Sharma, S.; Elson, M. J.; Spear, K. L.; Glidden, A. M.; Little, M. A.; et al. 2018. Using Smartphones and Machine Learning to Quantify Parkinson Disease Severity: The Mobile Parkinson Disease Score. *JAMA Neurology*.
- Zheng, B.; Yoon, S. W.; and Lam, S. S. 2014. Breast cancer diagnosis based on feature extraction using a hybrid of k-means and support vector machine algorithms. *Expert Systems with Applications* 41(4):1476–1482.