

A Deep Learning Approach to Diagnosing Multiple Sclerosis from Smartphone Data

Patrick Schwab, and Walter Karlen, *Senior Member, IEEE*

Abstract— Multiple sclerosis (MS) affects the central nervous system with a wide range of symptoms. MS can, for example, cause pain, changes in mood and fatigue, and may impair a person’s movement, speech and visual functions. Diagnosis of MS typically involves a combination of complex clinical assessments and tests to rule out other diseases with similar symptoms. New technologies, such as smartphone monitoring in free-living conditions, could potentially aid in objectively assessing the symptoms of MS by quantifying symptom presence and intensity over long periods of time. Here, we present a deep-learning approach to diagnosing MS from smartphone-derived digital biomarkers that uses a novel combination of a multilayer perceptron with neural soft attention to improve learning of patterns in long-term smartphone monitoring data. Using data from a cohort of 774 participants, we demonstrate that our deep-learning models are able to distinguish between people with and without MS with an area under the receiver operating characteristic curve of 0.88 (95% CI: 0.70, 0.88). Our experimental results indicate that digital biomarkers derived from smartphone data could in the future be used as additional diagnostic criteria for MS.

Index Terms— Artificial neural networks, digital biomarkers, medical diagnosis, multiple sclerosis, explainability

I. INTRODUCTION

MULTIPLE sclerosis (MS) is a neurological disease that affects around 2 million people worldwide [1]. The neural lesions caused by MS reduce the capability of neurons to transmit information, which leads to a wide range of symptoms, such as changes in sensation, mobility, balance, vision, and cognition [2]. Diagnosing MS requires objective evidence of two lesions in the central nervous system disseminated both in time and space [3], [4]. Physicians typically use a combination of clinical assessments of symptoms, blood tests, imaging, cerebrospinal fluid analysis and analysis of evoked potentials to rule out other diseases with similar symptoms [3], [4], [5]. Currently, no cure exists for MS, but there are treatments available that are effective at managing the symptoms of MS and may significantly improve long-term outcomes [2], [6], [7]. To receive early access to these treatments, a timely diagnosis is of paramount importance for patients.

This work was partially funded by the Swiss National Science Foundation (SNSF) project No. 167302 within NRP 75 “Big Data”. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPUs used for this research.

P. Schwab and W. Karlen are with the Mobile Health Systems Lab, Institute of Robotics and Intelligent Systems, Department of Health Sciences and Technology, ETH Zurich, Switzerland (e-mails: patrick.schwab@hest.ethz.ch, walter.karlen@ieee.org).

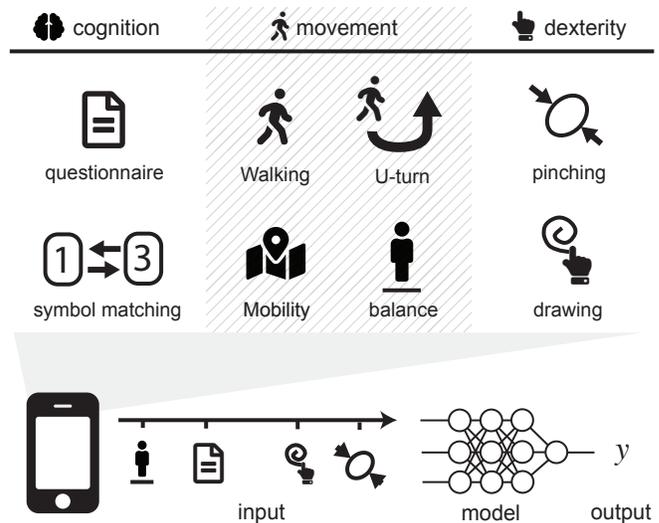


Fig. 1. Smartphone-based tests (top) can be used to assess cognition, movement and finger dexterity symptoms of multiple sclerosis (MS) and track their progression over time. We train machine-learning models to learn to produce a scalar diagnostic score y (bottom) from the data collected during any number of those tests to learn to diagnose MS.

Smartphone-based tests could potentially be used to quantify symptoms of MS in the wild over long periods of time (Figure 1). However, to date, it has not been established whether and to what degree smartphone monitoring data can be used to derive digital biomarkers for the diagnosis of MS. A particularly challenging aspect of using smartphone data to derive digital biomarkers for the diagnosis of MS is that smartphone monitoring yields large amounts of high-resolution data from multiple symptom categories. Identifying the salient input segments and reaching a clinically meaningful conclusion from raw sensor data in an accurate and timely manner is therefore challenging both for physicians and machines.

To address these issues, we present a machine-learning approach for distinguishing between people with and without MS from smartphone data. At its core, our method uses an attentive aggregation model (AAM) that integrates the results of multiple tests over time to produce a single diagnostic score. By integrating neural attention in our model, we are additionally able to quantify the importance of individual tests towards the model’s output. Our experiments on real-world smartphone monitoring data show that our method outperforms several strong baselines, identifies meaningful patterns, and that smartphone data could potentially be used to derive digital biomarkers for the diagnosis of MS.

Concretely, our contributions are as follows:

- We present a deep-learning approach to distinguishing between people with and without MS that integrates data from multiple types of smartphone-based tests performed over long time frames (up to more than 200 days).
- We extend our deep-learning approach with neural soft attention in order to quantify the importance of individual input features towards the final diagnostic score.
- We utilise real-world smartphone monitoring data from 774 subjects to evaluate, for the first time, the comparative performance of machine-learning models and several strong baselines in distinguishing between people with and without MS.

II. RELATED WORK

Using machine learning to aid in medical tasks has attracted much research interest. Researchers have, for example, used machine learning for mortality modelling [8], sepsis decision support [9], alarm reduction in critical care [10], to provide explainable decisions in medical decision-support systems [11], [12], [13], and to identify subtypes in autism spectrum disorders [14] and scleroderma [15]. Giving a reliable diagnosis is one of the most challenging tasks in medicine that requires significant domain knowledge in the disease being assessed, and the ability to integrate information from a large number of potentially noisy data sources. Machine learning is an attractive tool to perform automated diagnoses because it can draw upon the experience of millions of historical samples in its training data, and seamlessly integrate data from multiple disparate data sources. In previous studies, machine learning has, for example, been used to diagnose skin cancer from image data [16], glaucoma from standard automated perimetry [17], and a large number of diseases from electronic health records [18], [19] and lab test results [20]. However, obtaining objective data about symptoms and symptom progression over time is challenging in many diseases. For some diseases, wearable devices and smartphones have emerged as viable tools for gathering diagnostic data in the wild. Smartphones and telemonitoring have, for example, been used to gather digital biomarkers for the diagnosis of melanomas [21], bipolar disorder [22], cognitive function [23], and Parkinson’s disease [24], [25], [26], [27], [28]. The use of machine learning on high-resolution health data often requires specialised approaches to effectively deal with missingness [29], [30], long-term temporal dependencies [19], noise [31], heterogeneity [32], irregular sampling [18], sparsity [33], and multivariate input data [10], [28], [34]. In this work, we build on these advances to develop a novel approach to learning to diagnose MS from smartphone-derived digital biomarkers that addresses the aforementioned challenges.

A. Monitoring and Diagnosis of MS

The clinical state-of-the-art in monitoring symptoms and symptom progression in MS is based on a combination of clinical assessments, such as neurological exams, magnetic resonance imaging (MRI), and the Expanded Disability Status Scale (EDSS) [35], [36]. However, these tests can only be

performed at clinical centers by medical specialists. With dozens of mHealth apps available to manage MS on all major smartphone platforms, smartphone apps have recently emerged as a readily accessible alternative to non-invasively track symptoms of MS in the wild [37], [38]. Prior studies on the use mHealth in MS have, for example, evaluated telemedicine-enabled remote EDSS scoring [39], measurement devices for estimating walking ability [40] and fatigue [41], and machine learning for assessing gait impairment in MS [42]. Epidemiologically, demographic factors, such as age and sex, have been shown to be predictive of MS [43].

In contrast to existing works that focused on single daily-life aspects of already diagnosed MS patients, we present an approach to diagnosing MS from smartphone-derived digital biomarkers, and verify this approach on a real-world dataset collected from a MS cohort. Our machine-learning approach addresses multiple challenges in learning from sensor-based smartphone tests that are self-administered multiple times over long periods of time. Most notably, with the integration of a global neural soft attention mechanism, we enable the quantification of the importance of individual smartphone tests towards the final diagnostic score, overcome the challenges of missingness, sparse data, long-term temporal dependencies between tests, and multivariate data with irregular sampling.

III. METHODS AND MATERIALS

A. Smartphone Tests

We utilise data collected by the Floodlight Open study, a large smartphone-based observational study for MS [44], [45]. The de-identified dataset used in this work is openly available to researchers¹. In the study, participants were asked to actively perform a number of smartphone-based tests on a daily basis using their personal smartphones in the wild and without any clinical supervision (Figure 1). However, participants were free to choose when and if they performed the daily tests. Many participants therefore did not strictly adhere to the daily test protocol, and performed the tests irregularly. In addition to the manual tests, the app also passively collected movement data of the participants in order to determine their radius of living. The following tests were included in the study (Figure 1) [44]:

- ☑ **Mood Questionnaire.** In the mood questionnaire, participants were asked a single question about their current well-being. The answers were mapped to a scalar mood score that was recorded for each answer. The score was used to track changes of participants’ mood over time.
- ☑☑ **Symbol Matching.** In the symbol matching test, participants were presented with a mapping of symbols to numbers. Participants were then prompted with a single symbol from this mapping, and asked to translate the shown symbol into the corresponding number using an on-screen virtual keyboard. Once the user entered their response, a new symbol would be shown. The goal was to translate the presented symbols as quickly and accurately as possible in a fixed amount of time. As metrics, the average response time and the number of correct responses were recorded. There was also a baseline version

¹<https://floodlightopen.com>

of this test in which participants simply had to input the presented numbers directly without any intermediate mapping for which the same metrics were recorded.

- ✎ **Walking.** In the walking test, participants were asked to take a walk for two minutes, where ever they saw fit. Their smartphones recorded the number of steps taken during this walk in order to capture whether the participants' ability to walk was impaired.
- ↺ **U-turn.** In the U-turn test, participants were asked to walk as quickly as possible between two fixed points of their choice that should approximately have been four meters apart. Their smartphones recorded the number of turns and the average turn speed during this test in order to assess the participants' ability to turn.
- 🚶 **Balance.** In the balance test, participants were asked to stand still and hold their balance for a fixed amount of time. The app recorded their postural sway during this test in order to evaluate to what degree the participant was able to remain still. Impaired postural control and an increased risk of falling are symptoms commonly associated with MS [46].
- 📍 **Mobility.** The mobility test recorded the daily life space of the participant using their smartphone's location sensors. The mobility test was the only test that ran in the background and did not have to be manually activated.
- ✂ **Pinching.** In the pinching test, participants were presented with a series of virtual objects in varying locations on their smartphone screens. The participants were then asked to perform a pinching gesture using their fingers to squash the objects as quickly as possible. The app recorded the number of successfully squashed objects over a fixed amount of time, and which hand was used to perform the pinching gesture. The aim was to measure the participants' pinching ability which may be impaired in people with MS [47].
- 🎨 **Drawing.** In the drawing test, participants were asked to draw a sequence of shapes that was shown on their screens using their fingers - twice for each shape. The shapes represented a square, a circle, a figure eight, and a spiral. For each shape, the app recorded the best Hausdorff distance between the drawn and shown shapes.

B. Problem Statement

We consider the setting in which we are given a number k of tests, including test scores $s_i \in \mathbb{R}$, the time since the last test t_i in seconds, a one-hot encoded representation $m_i \in \{0, 1\}$ of the test metric type with $i \in [0 \dots k - 1]$, and demographic data $d = (d_0, d_1)$, including age $d_0 \in \mathbb{N}_0$ and sex $d_1 \in \{0, 1\}$, of the participant that performed the tests. Our goal is to train a predictive model P that produces a scalar diagnostic score $y \in [0, 1]$ that indicates the likelihood of the given set of test results belonging to a participant with or without MS.

$$y = P([0, t_1, \dots, t_i], [m_0, \dots, m_i], [s_0, \dots, s_i], d). \quad (1)$$

The primary challenge in this setting is to identify predictive patterns among the potentially large set of tests performed irregularly over long periods of time that provide evidence for or against a MS diagnosis.

C. Attentive Aggregation Model (AAM)

As predictive model P , we use an AAM - a deep-learning model that utilises neural soft attention in order to integrate information from a potentially large number of smartphone test results. AAMs are based on evidence aggregation models (EAMs) [28]. As a first step, we concatenate the time since the respective prior test t_i , result scores s_i , and test metric indicators m_i from the smartphone tests into k features x_i .

$$x_i = \text{concatenate}(t_i, m_i, s_i) \quad (2)$$

We then use a multilayer perceptron (MLP) with a configurable number L of hidden layers and N neurons per hidden layer to process each input feature into a N -dimensional high-level hidden feature representation h_i .

$$h_i = \text{MLP}(x_i) \quad (3)$$

Next, we aggregate the information from all k high-level hidden representations into a single aggregated hidden representation h_{all} that reflects all available tests for a given participant. To do so, we used a learned neural soft attention mechanism that weighs the individual hidden representations h_i of each individual test instance by their respective importance a_i towards the final diagnostic output score.

$$h_{\text{all}} = \sum_{i=1}^k a_i h_i \quad (4)$$

Following [13], [31], [48], we calculate the attention factors a_i by first projecting the individual hidden representations h_i into an attention space representation u_i through a single-layer MLP with a weight matrix W and bias b .

$$u_i = \text{activation}(W h_i + b) \quad (5)$$

We finalise the calculation of the attention factors a_i by computing a softmax over the attention space representations u_i of h_i using the most informative hidden representation u_{max} . W , b and u_{max} are learned parameters and optimised together with the other parameters during training [28].

$$a_i = \text{softmax}(u_i^T u_{\text{max}}) \quad (6)$$

We then calculate the final diagnostic score y using a MLP with one sigmoid output node on a concatenation of the aggregated hidden state h_{all} and the demographic data d . The integration of demographic information to the top-level MLP enables the model to adjust the final prediction y for epidemiological risk factors [43], such as sex and age.

$$y = \text{MLP}(\text{concatenate}(h_{\text{all}}, d)) \quad (7)$$

AAMs rely solely on attention to aggregate test results over time. Using attention to perform the temporal aggregation has the advantage that attention mechanisms learn global input-output relations without regard to their distance in the input sequence, and in this manner improve learning of long-range temporal dependencies [49]. While global attention models are commonly employed in natural language processing [49], [50], we are not aware of any prior works that apply global attention to improve learning of long-term temporal dependencies on smartphone sensor data.

TABLE I

POPULATION STATISTICS. Training, validation, and test fold statistics. Age and Usage are medians (10% and 90% quantiles in parentheses).

Property	Training	Validation	Test
Subjects (#)	542 (70%)	77 (10%)	155 (20%)
MS (%)	51.9	52.0	51.6
Female (%)	60.3	59.7	60.7
Age (years)	41.0 (27.0, 59.0)	41.0 (26.5, 56.5)	41.0 (28.0, 57.5)
Usage (days)	22.4 (0.6, 203.5)	19.0 (0.8, 175.5)	18.8 (1.0, 130.0)

IV. EXPERIMENTS

To evaluate the predictive performance of AAMs in diagnosing MS from smartphone data, we performed experiments that compared the diagnostic performance of AAMs and several baseline models on real-world smartphone monitoring data. Our experiments aimed to answer the following questions:

- 1 What is the predictive performance of AAMs in diagnosing MS from smartphone data?
- 2 What is the predictive performance of AAMs compared to other methods, such as Mean Aggregation and the demographic baseline?
- 3 Which test types were most informative for diagnosing MS, and to what degree?
- 4 To what degree does including more tests performed by subjects improve predictive performance of AAMs?
- 5 Does the neural attention mechanism identify meaningful patterns?

A. Dataset and Study Cohort

We used data from the Floodlight Open study that recruited participants via their personal smartphones in, among others, the United States (US), Canada, Denmark, Spain, Italy, the Czech Republic, and Switzerland [44], [45]. To perform our experimental comparison, we used all of the available smartphone monitoring data from April 23rd 2018 to August 29th 2019. In addition to regularly performing the smartphone-based tests, participants also provided their demographic profiles upon sign-up. The demographic profile included age, sex, and whether or not they had an existing diagnosis of MS. To ensure that a minimal amount of data points are available for diagnosis, we excluded all participants that had produced fewer than 20 test results during the analysed time frame (312 participants). We chose the cutoff at a minimum of 20 tests as this corresponds roughly to two sets of the daily test suite, which would be the minimal amount of data needed to assess symptom progression over time. We assigned the included patients to three folds for training (70%), validation (10%) and testing (20%) randomised within strata of diagnostic status, app usage, sex, age, and number of tasks performed (Table I). The training, validation and test folds were respectively solely used for training, hyperparameter optimisation, and model evaluation. On average, participants with MS performed more tests over their app usage duration than participants without MS - presumably because of their intrinsic motivation to better understand their disease and help advance biomedical research.

B. Models

We trained a demographic baseline model, a Mean Aggregation baseline, and several ablations of AAMs. The AAMs used a fully-connected neural network as their base, and a single neuron with a sigmoid activation function as output. We trained one AAM version that received the demographic information (AAM + age + sex), and one that did not (AAM). For computational reasons, we limited the maximum number of test results per participant. To estimate the performance benefit of having access to information from more tests in the analysis, we trained AAMs and Mean Aggregation using up to the first 25, 30, 40, 50, 100, 150, 200, 250, 300 and 350 test results per participant, if available. For the demographic baseline, we used a random forest (RF) model that received as input only the age and sex of the participant (Age + sex). We used the demographic baseline to evaluate whether and to what degree data from the smartphone-based tests improves predictive performance, since the demographic baseline only had access to demographic data and did not include data from the smartphone-based tests. As a simple reference baseline, the Mean Aggregation utilised the mean normalised test result score to produce the final diagnostic score - this served to determine whether the use of more complex, learned aggregation methods, such as AAMs, is effective and warranted.

C. Hyperparameters

To ensure all models were given the same degree of hyperparameter optimisation, we used a standardised approach where each model was given an optimisation budget of 50 hyperparameter optimisation runs with hyperparameters chosen from pre-defined ranges (Table II). The best performing configuration on the validation fold was selected for further evaluation. For the demographic baseline model, we used a RF with T trees and a maximum tree depth of D . For the AAMs, we used an initial MLP with L hidden layers, N hidden units per hidden layer, a dropout percentage of p between hidden layers, and an L2 weight penalty of strength s . We trained AAMs to optimise binary cross entropy for a maximum of 300 epochs with a minibatch size of B participants and a learning rate of 0.003. In addition, we used early stopping with a patience of 32 epochs on the validation fold.

TABLE II

HYPERPARAMETERS. Ranges used for hyperparameter optimisation of AAMs (top) and the Age + sex baseline using a Random Forest model (RF, bottom). Parentheses indicate continuous ranges within the indicated limits sampled at uniform probability. Comma-delimited lists indicate discrete choices with equal selection probability.

	Hyperparameter	Range / Choices
AAM	Number of hidden units N	16, 32, 64, 128
	Batch size B	16, 32, 64
	L2 regularisation strength s	0.0001, 0.00001, 0.0
	Number of layers L	(1, 3)
	Dropout percentage p	(0%, 35%)
RF	Tree depth D	3, 4, 5
	Number of trees T	32, 64, 128, 256

TABLE III

PREDICTIVE PERFORMANCE. Comparison of Attentive Aggregation Models (AAMs), Mean Aggregation, and a demographic baseline (Age + sex) in terms of AUC, AUPR, F_1 , sensitivity, and specificity for predicting MS on the test set using a maximum of 250 test results from each participant. In parentheses are the 95% confidence intervals (CIs) obtained via bootstrap resampling. † = significant at $p < 0.05$ to AAM + age + sex.

Model (max. 250 test results)	AUC	AUPR	F_1	Sensitivity	Specificity
AAM + age + sex	0.88 (0.70, 0.88)	0.90 (0.67, 0.90)	0.80 (0.65, 0.83)	0.83 (0.59, 0.86)	0.73 (0.62, 0.89)
Mean Aggregation + age + sex	†0.77 (0.70, 0.82)	†0.76 (0.64, 0.84)	†0.71 (0.65, 0.78)	†0.68 (0.59, 0.85)	†0.75 (0.59, 0.87)
Age + sex	†0.76 (0.69, 0.84)	†0.75 (0.65, 0.86)	†0.69 (0.62, 0.79)	†0.73 (0.55, 0.83)	†0.61 (0.58, 0.89)
AAM	†0.72 (0.56, 0.82)	†0.67 (0.57, 0.84)	†0.61 (0.53, 0.77)	†0.63 (0.45, 0.79)	†0.83 (0.53, 0.86)
Mean Aggregation	†0.56 (0.50, 0.67)	†0.61 (0.49, 0.74)	†0.39 (0.20, 0.54)	†0.28 (0.13, 0.43)	†0.85 (0.70, 0.97)

D. Preprocessing

We normalised the time between two test results t_i to the range $[0, 1]$ using the highest observed t_i on the training fold of the dataset. We additionally normalised all test result scores x_i to the range of $[0, 1]$ using the lowest and highest observed test result for each test metric on the training fold of the dataset.

E. Metrics

1) *Predictive Performance*: We evaluated all models in terms of their area under the receiver operating characteristic curve (AUC), the area under the precision recall curve (AUPR), and F_1 score on the test fold of 155 participants. For the comparison of predictive performance, we additionally computed the sensitivity and specificity of the respective models. We also quantified the uncertainty of all the performance metrics by computing 95% confidence intervals (CIs) using bootstrap resampling with 1000 bootstrap samples. To assess the statistical significance of our results, we applied Mann-Whitney-Wilcoxon (MWW) tests at significance level $\alpha = 0.05$ to the main comparisons. We additionally applied the Bonferroni correction to adjust for multiple comparisons.

2) *Importance of Test Types*: To quantify the importance of the various test types toward the diagnostic performance of the AAM, we retrained AAMs with the same hyperparameters after removing the test results from exactly one type of test. The reduction in predictive performance associated with removing the information of one test type can be seen as a proxy for the importance of that test type [13], [51], since features that

are associated with a higher reduction in prediction error carry more weight in improving the model’s ability to predict MS.

3) *Neural Attention*: In order to qualitatively inspect the patterns that were captured by the AAM in the data, we additionally plotted the attention assigned to the test results from a sample participant with MS over time.

V. RESULTS

A. Predictive Performance

In terms of predictive performance for diagnosing MS, the AAMs models with demographic information (AAM + age + sex) achieved a significantly ($p < 0.05$) higher AUC, AUPR and F_1 than all the baselines we evaluated (Table III). In particular, we found that integrating the information from the smartphone tests was crucial, as AAMs that used both the demographic data and the smartphone test information significantly ($p < 0.05$) outperformed the demographic baseline (Age + sex) in terms of AUC, AUPR, and F_1 . Mean Aggregation had a considerably lower performance than AAMs, and displayed the worst AUC, AUPR and F_1 of the compared models - demonstrating that the use of more sophisticated adaptive aggregation models for integrating information from smartphone-based tests over time, such as AAMs, is effective and warranted. We also found that AAMs achieved a high level of both sensitivity and specificity, whereas the demographic model emphasised sensitivity, and Mean Aggregation specificity. As expected, the AAM achieved a significantly higher performance when it had access to demographic information, as it was able to adjust for demographic risk factors by, for

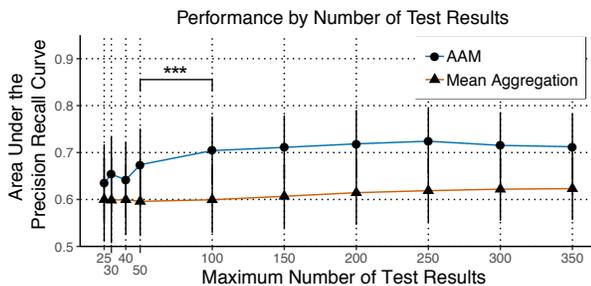


Fig. 2. Performance comparison of AAM (dots, blue) and Mean Aggregation (triangles, orange) in terms of their Area Under the Precision-Recall Curve (AUPR, y-axis) when varying the maximum number of test results (x-axis) available to predict the MS diagnosis for each participant. *** = significant at $p < 0.001$.

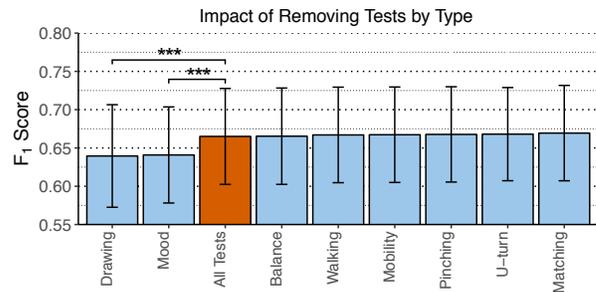


Fig. 3. Performance comparison of AAM in terms of their F_1 score (y-axis) in predicting the MS diagnosis for each participant after removal of the information of all tests of a specific type (labelled test types, bottom) from the dataset. The reference baseline without removal of any test types (All Tests) is highlighted in orange. *** = significant at $p < 0.001$.

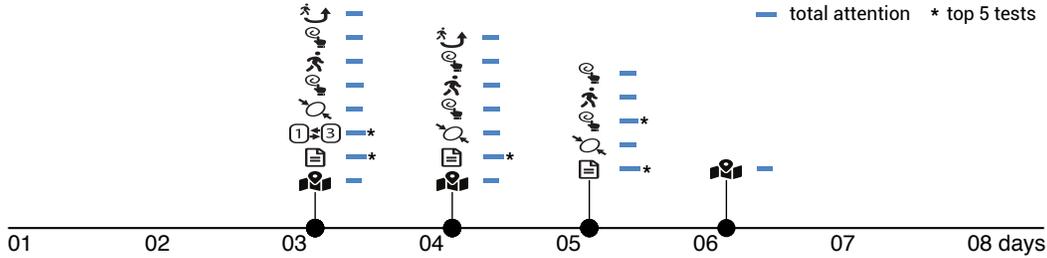


Fig. 4. A set of test results as performed by a female participant (FL51683656) aged 49 with MS from the test set. The timeline depicts usage days from earlier (left) to later (right). The black dots indicate that at least one test was performed that day. The absence of a dot indicates that no test was performed that day. Symbols connected to marked days depict which tests were performed (Figure 1). The blue bars directly adjacent to the symbols show how much total attention was assigned to that test instance. Total attention refers to the sum of attention assigned to all test results belonging to a test instance, e.g. the Hausdorff distances for all shapes in a drawing test. We found that the model focused predominantly on the mood, symbol matching, and drawing tests (* marks top 5 tests) to correctly identify (score = 0.77, threshold = 0.49) that this participant has MS.

example, assigning higher risk to female app users - which have an up to three times higher MS risk than men [43].

B. Impact of Performing More Tests

The results from comparing AAMs and Mean Aggregation in terms of their AUPR across a wide range of numbers of test results indicated that AAMs are better able to leverage an increasing number of tests performed (Figure 2). In terms of predictive performance as measured by the AUPR, the AAM surpassed the Mean Aggregation baseline at all evaluated maximum numbers of test results. We also found that having access to a higher number of tests consistently and significantly (50 vs. 100 test results, $p < 0.001$) improved the performance of AAMs up to a maximum of 250 test results per participant. Our results indicate that a respectable level of predictive performance could be achieved after collecting just 100 test results. In practice, data collection could for example be initiated upon referral to a specialist clinic or between the first clinical contact and the diagnostic appointment. We note that the time required to collect 100 test results is small in relation to the median time between symptom onset and diagnosis, which may be over two years in some countries [52]. In addition, Floodlight Open tests have been reported to be well accepted by patients with good adherence and patient satisfaction over a monitoring period of 24 weeks [45].

C. Importance of Test Types

When comparing the marginal reduction in prediction error associated with removing a specific type of test from the set of available tests, we found that the drawing and mood tests were contributing significantly larger marginal reductions ($p < 0.001$) in prediction error to the AAM (Figure 3) - indicating that the drawing and mood tests were more predictive of MS diagnosis than other test types. The removal of other test types did not lead to similarly considerable reductions in prediction error compared to the AAM that had access to information from all test types (All Tests). This result could indicate that the results of other test types were either (i) not highly predictive of MS, or (ii) correlated with other tests to a degree that strongly impacted their marginal contributions. Further prospective studies will be necessary to determine

the optimal set of tests because confounding factors, such as different usage patterns in response to different test suites being available, may also influence predictive performance.

D. Neural Attention

Qualitatively, on the data from one sample participant, we found that the model focused on mood, drawing, and symbol matching tests to diagnose MS (Figure 4). The focus on mood and drawing tests for this subject are in line with our findings on the overall importance of the test types (Table 3).

VI. DISCUSSION

To the best of our knowledge, this work is the first to present a machine-learning approach to diagnosing MS from long-term smartphone monitoring data collected outside of the clinical environment. To derive a scalar diagnostic score for MS from smartphone monitoring data, we used an AAM to aggregate data from multiple types of tests over long periods of time. The AAMs used neural attention to quantify the degree to which individual tests contribute to the final diagnostic score, and to overcome the challenges of missingness, sparse data, long-term temporal dependencies and irregular sampling. Our experimental results indicate that AAMs outperform several strong baselines, identify meaningful data points within the set of performed tests, and that smartphone-based digital biomarkers could potentially be used to aid in diagnosing MS alongside existing clinical tests. Among the several potential advantages of using smartphone-derived digital biomarkers for diagnosing MS are that smartphone-based tests (i) can be administered remotely and therefore potentially expand access to underserved geographic regions, (ii) are inexpensive to distribute and could therefore potentially become a low-cost alternative to more expensive in-clinic tests, and (iii) are able to integrate information from long-term symptom monitoring, and therefore potentially better represent and quantify fluctuations in symptom burden over time. An additional benefit of using smartphone-based diagnostics for MS is that machine-learning models can, as demonstrated in this work, identify which symptom categories are most indicative of MS and how they interact over time, and could therefore potentially be used to monitor disease progression and inform follow-up treatment decisions. Finally, our results show that information from

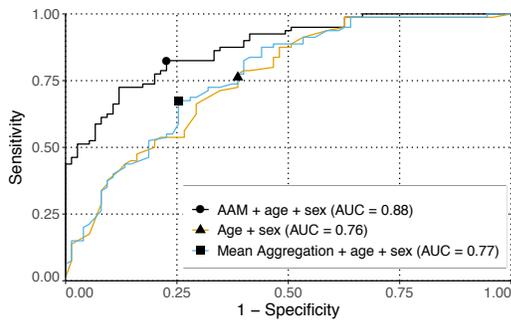


Fig. 5. Comparison of the receiver operating characteristic (ROC) curves of AAM + age + sex (black, dot), Age + sex (orange, triangle), and Mean Aggregation + age + sex (blue, square) when using a maximum of 250 test results to predict the MS diagnosis for each test set participant. Symbols indicate the operating points (thresholds selected on the validation fold) presented in the comparison in Table III.

smartphone-based biomarkers are to some degree orthogonal to more traditional measurements, such as demographic data, and could therefore potentially be integrated with information from existing clinical tests and other multimodal data sources, such as MRI, to further increase diagnostic accuracy.

A. Limitations

While of respectable size, the studied cohort was restricted to residents of a limited number of countries and likely not representative of the global population. More importantly, the data originated from patients with a diagnosis at various stages of the disease, instead of pre-diagnosis. Therefore, our work cannot conclude whether people without clinical MS diagnosis could be identified before they receive their diagnosis. A prospective validation in a larger, clinically representative cohort will be necessary to conclusively establish the performance and utility of smartphone-derived biomarkers as a tool to aid in the diagnosis of MS, and their robustness when confronted with other disorders that have similar symptoms. Further work should investigate whether such biomarkers are also suitable to track disease progression, identify and predict relapses, enable the effective tuning of therapeutic options and medication dosages, enable earlier diagnoses, and eventually also provide an accurate prognosis.

VII. CONCLUSION

We presented a novel machine-learning approach to distinguishing between people with and without MS based on long-term smartphone monitoring data. Our method uses an AAM to aggregate data from multiple types of tests over long periods of time to produce a scalar diagnostic score. AAMs use neural attention to quantify the degree to which individual tests contributed to the diagnostic score, and to overcome the challenges of missingness, sparse data, long-term temporal dependencies and irregular sampling. In an experimental evaluation on smartphone monitoring data from a cohort of 774 people, we demonstrated that AAMs identify predictive and meaningful digital biomarkers for diagnosing MS. Our experiments show that smartphone-derived digital biomarkers

could potentially be used to aid in diagnosing MS in the future alongside existing clinical tests. Smartphone-based tools for tracking symptoms and symptom progression in MS may improve clinical decision-making by giving clinicians access to objectively measured high-resolution health data collected outside the restricted clinical environment. In addition, our solution based on attention mechanisms further elucidates the basis of the model decisions and may enhance the clinician's understanding of the provided diagnostic score. We therefore believe our initial results may warrant further research on how digital biomarkers could be integrated into clinical workflows.

ACKNOWLEDGMENTS

The data used in this manuscript were contributed by users of the Floodlight Open mobile application developed by Genentech Inc: <https://floodlightopen.com>. Patrick Schwab is an affiliated PhD fellow at the Max Planck ETH Center for Learning Systems.

REFERENCES

- [1] T. Vos, C. Allen, M. Arora, R. M. Barber, Z. A. Bhutta, A. Brown, A. Carter, D. C. Casey, F. J. Charlson, A. Z. Chen *et al.*, "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*, vol. 388, no. 10053, pp. 1545–1602, 2016.
- [2] W. J. Brownlee, T. A. Hardy, F. Fazekas, and D. H. Miller, "Diagnosis of multiple sclerosis: Progress and challenges," *The Lancet*, vol. 389, no. 10076, pp. 1336–1346, 2017.
- [3] W. I. McDonald, A. Compston, G. Edan, D. Goodkin, H.-P. Hartung, F. D. Lublin, H. F. McFarland, D. W. Paty, C. H. Polman, S. C. Reingold *et al.*, "Recommended diagnostic criteria for multiple sclerosis: Guidelines from the International Panel on the diagnosis of multiple sclerosis," *Annals of Neurology: Official Journal of the American Neurological Association and the Child Neurology Society*, vol. 50, no. 1, pp. 121–127, 2001.
- [4] C. H. Polman, S. C. Reingold, B. Banwell, M. Clanet, J. A. Cohen, M. Filippi, K. Fujihara, E. Havrdova, M. Hutchinson, L. Kappos *et al.*, "Diagnostic criteria for multiple sclerosis: 2010 revisions to the McDonald criteria," *Annals of Neurology*, vol. 69, no. 2, pp. 292–302, 2011.
- [5] M. Filippi, M. A. Rocca, O. Ciccarelli, N. De Stefano, N. Evangelou, L. Kappos, A. Rovira, J. Sastre-Garriga, M. Tintorè, J. L. Frederiksen *et al.*, "MRI criteria for the diagnosis of multiple sclerosis: MAGNIMS consensus guidelines," *The Lancet Neurology*, vol. 15, no. 3, pp. 292–303, 2016.
- [6] M. Marziniak, K. Ghorab, W. Kozubski, C. Pflieger, L. Sousa, K. Vernon, M. Zaffaroni, and S. G. Meuth, "Variations in multiple sclerosis practice within Europe – Is it time for a new treatment guideline?" *Multiple Sclerosis and Related Disorders*, vol. 8, pp. 35–44, 2016.
- [7] N. Scolding, D. Barnes, S. Cader, J. Chataway, A. Chaudhuri, A. Coles, G. Giovannoni, D. Miller, W. Rashid, K. Schmierer *et al.*, "Association of British Neurologists: Revised (2015) guidelines for prescribing disease-modifying treatments in multiple sclerosis," *Practical Neurology*, vol. 15, no. 4, pp. 273–279, 2015.
- [8] M. Ghassemi, T. Naumann, F. Doshi-Velez, N. Brimmer, R. Joshi, A. Rumshisky, and P. Szolovits, "Unfolding physiological state: Mortality modelling in intensive care units," in *SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2014, pp. 75–84.
- [9] S. Horng, D. A. Sontag, Y. Halpern, Y. Jernite, N. I. Shapiro, and L. A. Nathanson, "Creating an automated trigger for sepsis clinical decision support at emergency department triage using machine learning," *PLoS one*, vol. 12, no. 4, p. e0174708, 2017.
- [10] P. Schwab, E. Keller, C. Muroi, D. J. Mack, C. Strässle, and W. Karlen, "Not to cry wolf: Distantly supervised multitask learning in critical care," in *International Conference on Machine Learning*, 2018.
- [11] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems*, 2017, pp. 4768–4777.

- [12] S. M. Lundberg, B. Nair, M. S. Vavilala, M. Horibe, M. J. Eisses, T. Adams, D. E. Liston, D. K.-W. Low, S.-F. Newman, J. Kim, and S.-I. Lee, "Explainable machine-learning predictions for the prevention of hypoxaemia during surgery," *Nature Biomedical Engineering*, vol. 2, no. 10, p. 749, 2018.
- [13] P. Schwab, D. Miladinovic, and W. Karlen, "Granger-causal attentive mixtures of experts: Learning important features with neural networks," in *AAAI Conference on Artificial Intelligence*, 2019.
- [14] F. Doshi-Velez, Y. Ge, and I. Kohane, "Comorbidity clusters in autism spectrum disorders: An electronic health record time-series analysis," *Pediatrics*, vol. 133, no. 1, pp. e54–e63, 2014.
- [15] P. Schulam, F. Wigley, and S. Saria, "Clustering longitudinal clinical marker trajectories from electronic health data: Applications to phenotyping and endotype discovery," in *AAAI Conference on Artificial Intelligence*, 2015.
- [16] A. Esteva, B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau, and S. Thrun, "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, p. 115, 2017.
- [17] K. Chan, T.-W. Lee, P. A. Sample, M. H. Goldbaum, R. N. Weinreb, and T. J. Sejnowski, "Comparison of machine learning and traditional classifiers in glaucoma diagnosis," *IEEE Transactions on Biomedical Engineering*, vol. 49, no. 9, pp. 963–974, 2002.
- [18] Z. C. Lipton, D. C. Kale, C. Elkan, and R. Wetzl, "Learning to diagnose with LSTM recurrent neural networks," in *International Conference on Learning Representations*, 2015.
- [19] E. Choi, M. T. Bahadori, A. Schuetz, W. F. Stewart, and J. Sun, "Doctor AI: Predicting clinical events via recurrent neural networks," in *Machine Learning for Healthcare Conference*, 2016.
- [20] N. Razavian, J. Marcus, and D. Sontag, "Multi-task prediction of disease onsets from longitudinal laboratory tests," in *Machine Learning for Healthcare Conference*, 2016.
- [21] J. A. Wolf, J. F. Moreau, O. Akilov, T. Patton, J. C. English, J. Ho, and L. K. Ferris, "Diagnostic inaccuracy of smartphone applications for melanoma detection," *JAMA Dermatology*, vol. 149, no. 4, 2013.
- [22] M. Faurholt-Jepsen, J. Busk, H. Þórarinsdóttir, M. Frost, J. E. Bardram, M. Vinberg, and L. V. Kessing, "Objective smartphone data as a potential diagnostic marker of bipolar disorder," *Australian & New Zealand Journal of Psychiatry*, vol. 53, no. 2, pp. 119–128, 2019.
- [23] A. Piau, K. Wild, N. Mattek, and J. Kaye, "Current state of digital biomarker technologies for real-life, home-based monitoring of cognitive function for mild cognitive impairment to mild Alzheimer disease and implications for clinical care: Systematic review," *Journal of Medical Internet Research*, vol. 21, no. 8, p. e12785, 2019.
- [24] A. Tsanas, M. A. Little, P. E. McSharry, and L. O. Ramig, "Enhanced classical dysphonia measures and sparse regression for telemonitoring of Parkinson's disease progression," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2010, pp. 594–597.
- [25] S. Arora, V. Venkataraman, A. Zhan, S. Donohue, K. Biglan, E. Dorsey, and M. Little, "Detecting and monitoring the symptoms of Parkinson's disease using smartphones: A pilot study," *Parkinsonism & Related Disorders*, vol. 21, no. 6, pp. 650–653, 2015.
- [26] A. Tsanas, M. A. Little, P. E. McSharry, and L. O. Ramig, "Accurate telemonitoring of Parkinson's disease progression by noninvasive speech tests," *IEEE Transactions on Biomedical Engineering*, vol. 57, no. 4, pp. 884–893, 2010.
- [27] A. Zhan, S. Mohan, C. Tarolli, R. B. Schneider, J. L. Adams, S. Sharma, M. J. Elson, K. L. Spear, A. M. Glidden, M. A. Little *et al.*, "Using smartphones and machine learning to quantify Parkinson disease severity: The Mobile Parkinson Disease Score," *JAMA Neurology*, 2018.
- [28] P. Schwab and W. Karlen, "PhoneMD: Learning to diagnose Parkinson's disease from smartphone data," in *AAAI Conference on Artificial Intelligence*, 2019.
- [29] Z. C. Lipton, D. C. Kale, and R. Wetzl, "Directly modeling missing data in sequences with RNNs: Improved classification of clinical time series," in *Machine Learning for Healthcare Conference*, 2016.
- [30] Z. Che, S. Purushotham, K. Cho, D. Sontag, and Y. Liu, "Recurrent neural networks for multivariate time series with missing values," *Scientific reports*, vol. 8, no. 1, p. 6085, 2018.
- [31] P. Schwab, G. C. Scabbia, J. Zhang, M. Delai, and W. Karlen, "Beat by beat: Classifying cardiac arrhythmias with recurrent neural networks," in *Computing in Cardiology*, 2017.
- [32] M. W. Libbrecht and W. S. Noble, "Machine learning applications in genetics and genomics," *Nature Reviews Genetics*, vol. 16, no. 6, pp. 321–332, 2015.
- [33] T. A. Lasko, J. C. Denny, and M. A. Levy, "Computational phenotype discovery using unsupervised feature learning over noisy, sparse, and irregular clinical data," *PLoS one*, vol. 8, no. 6, p. e66341, 2013.
- [34] M. Ghassemi, M. A. Pimentel, T. Naumann, T. Brennan, D. A. Clifton, P. Szolovits, and M. Feng, "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*, 2015.
- [35] J. F. Kurtzke, "Rating neurologic impairment in multiple sclerosis: An expanded disability status scale (EDSS)," *Neurology*, vol. 33, no. 11, pp. 1444–1452, 1983.
- [36] M. P. Wattjes, À. Rovira, D. Miller, T. A. Yousry, M. P. Sormani, N. De Stefano, M. Tintoré, C. Auger, C. Tur, M. Filippi *et al.*, "Evidence-based guidelines: MAGNIMS consensus guidelines on the use of MRI in multiple sclerosis - establishing disease prognosis and monitoring patients," *Nature Reviews Neurology*, vol. 11, no. 10, p. 597, 2015.
- [37] G. Giunti, E. G. Fernández, E. D. Zubiete, and O. R. Romero, "Supply and demand in mHealth apps for persons with multiple sclerosis: Systematic search in app stores and scoping literature review," *JMIR mHealth and uHealth*, vol. 6, no. 5, 2018.
- [38] A. K. Boukhalova, O. Fan, A. M. Weideman, T. Harris, E. Kowalczyk, P. Kosa, and B. Bielekova, "Smartphone level test measures disability in several neurological domains for patients with multiple sclerosis," *Frontiers in Neurology*, vol. 10, p. 358, 2019.
- [39] R. Bove, C. Bevan, E. Crabtree, C. Zhao, R. Gomez, P. Garcha, J. Morrissey, J. Dierkhising, A. J. Green, S. L. Hauser *et al.*, "Toward a low-cost, in-home, telemedicine-enabled assessment of disability in multiple sclerosis," *Multiple Sclerosis Journal*, vol. 25, no. 11, pp. 1526–1534, 2018.
- [40] G. Dalla-Costa, M. Radaelli, S. Maida, F. Sangalli, B. Colombo, L. Moiola, G. Comi, and V. Martinelli, "Smart watch, smarter EDSS: Improving disability assessment in multiple sclerosis clinical practice," *Journal of the Neurological Sciences*, vol. 383, pp. 166–168, 2017.
- [41] L. Barrios, P. Oldrati, S. Santini, and A. Lutterotti, "Recognizing digital biomarkers for fatigue assessment in patients with multiple sclerosis," in *EAI International Conference on Pervasive Computing Technologies for Healthcare – Demos, Posters, Doctoral Colloquium*, 8 2018.
- [42] R. S. McGinnis, N. Mahadevan, Y. Moon, K. Seagers, N. Sheth, J. A. Wright Jr, S. DiCristofaro, I. Silva, E. Jortberg, M. Ceruolo *et al.*, "A machine learning approach for gait speed estimation using skin-mounted wearable sensors: From healthy controls to individuals with multiple sclerosis," *PLoS one*, vol. 12, no. 6, p. e0178366, 2017.
- [43] M. Pugliatti, G. Rosati, H. Carton, T. Riise, J. Drulovic, L. Vécsei, and I. Milanov, "The epidemiology of multiple sclerosis in Europe," *European Journal of Neurology*, vol. 13, no. 7, pp. 700–722, 2006.
- [44] X. Montalban, P. Mulero, L. Midaglia, J. Graves, S. Hauser, L. Julian, M. Baker, J. Schadrack, C. Gossens, A. Scotland *et al.*, "FLOODLIGHT: Remote self-monitoring is accepted by patients and provides meaningful, continuous sensor-based outcomes consistent with and augmenting conventional in-clinic measures," *Neurology 92 (Supplement)*, 2018.
- [45] L. Midaglia, P. Mulero, X. Montalban *et al.*, "Adherence and satisfaction of smartphone- and smartwatch-based remote active testing and passive monitoring in people with multiple sclerosis: Nonrandomized interventional feasibility study," *JMIR*, vol. 21, no. 8, 2019.
- [46] M. H. Cameron and S. Lord, "Postural control in multiple sclerosis: implications for fall prevention," *Current Neurology and Neuroscience Reports*, vol. 10, no. 5, pp. 407–412, 2010.
- [47] C. C. Chen, N. Kasven, H. I. Karpatkin, and A. Sylvester, "Hand strength and perceived manual ability among patients with multiple sclerosis," *Archives of Physical Medicine and Rehabilitation*, vol. 88, no. 6, pp. 794–797, 2007.
- [48] K. Xu, J. Ba, R. Kiro, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio, "Show, attend and tell: Neural image caption generation with visual attention," in *International Conference on Machine Learning*, 2015, pp. 2048–2057.
- [49] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, E. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, 2017, pp. 5998–6008.
- [50] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [51] P. Schwab and W. Karlen, "CXPlain: Causal Explanations for Model Interpretation under Uncertainty," in *Advances in Neural Information Processing Systems*, 2019.
- [52] O. Fernández, V. Fernández, T. Arbizu, G. Izquierdo, I. Bosca, R. Arroyo, J. G. Merino, E. De Ramón *et al.*, "Characteristics of multiple sclerosis at onset and delay of diagnosis and treatment in Spain (The Novo Study)," *Journal of Neurology*, vol. 257, no. 9, pp. 1500–1507, 2010.