ANNALS of Clinical and Translational Neurology
Open Access
AMERICAN NEUROLOGICAL ASSOCIATION

RESEARCH ARTICLE

A machine learning algorithm successfully screens for Parkinson's in web users

Brit Youngmann[1], Liron Allerhand[1], Ora Paltiel[2], Elad Yom-Tov[1,3] & David Arkadir[4]

[1]Microsoft Research, Herzliya, Israel
[2]Braun School of Public Health and Community Medicine, Hadassah-Hebrew University, Jerusalem, Israel
[3]Faculty of Industrial Engineering and Management, Technion, Haifa, Israel
[4]Department of Neurology, Hadassah Hebrew University, Jerusalem, Israel

Correspondence
David Arkadir, Hadassah Medical Center, POB 12000, Jerusalem 91120, Israel.
Tel: +972 2 6777111;
E-mail: arkadir@hadassah.org.il

Funding Information
No funding information provided.

Received: 10 October 2019; Accepted: 21 October 2019

Annals of Clinical and Translational Neurology **2019; 6(12): 2503–2509**

doi: 10.1002/acn3.50945

Abstract

Objective: To develop, apply, and evaluate, a novel web-based classifier for screening for Parkinson disease among a large cohort of search engine users. **Methods**: A supervised machine learning classifier learned to distinguish web users with self-reported Parkinson's disease from controls based on their interactions with a search engine (Bing, Microsoft). It was then applied to groups of web users with low or high risk for actual Parkinson's disease. Textual content of web queries was used to sort surfers into the different risk groups, but not for classifying users as negative or positive for Parkinson's disease. Disease detection was unsolicited. Researchers did not have access to any identifying data on users. **Results**: Applying the classifier (with an estimated positive predictive value of 25%) resulted in 17,843/1,490,987 (1.2%) web users over the age of 40 years screened positive for Parkinson's disease. This percentile was higher in at-risk groups (Fisher exact $P < 0.00001$), including users who searched for information regarding the disease (518/804, 64.4%), and users with non-motor Parkinson's symptom or with an affected relative (57/1064, 5.3%). Longitudinal follow-up revealed that in all studied groups individuals classified as having the disease showed a higher mean rate of progression in disease-related features (t-test $P < 0.05$). **Interpretation**: An automatic classifier, based on mouse and keyboard interactions with a search engine, is able to reliably trace individuals at high risk for actual Parkinson's disease as well as to demonstrate more rapid progression of disease-related signs in those who screened positive. This ability raises novel ethical issues.

Introduction

Massive volumes of data, collected on the web from millions of users, can be analyzed to reveal disease trends in the population. Studies have employed computer algorithms to detect[1,2] and predict[3] infectious disease epidemics in real time, track health behaviors,[4] and assess the post-marketing safety profiles of drugs.[5] Recent refinement of machine learning algorithms has expanded the focus from epidemiology to the health conditions of individual web users. Algorithms were developed for use with data from search engine queries to identify web users with such diseases as cancer[6,7] or Parkinson's disease (PD).[8,9] However, such tools have remained untested and no attempt has been made to validate their diagnostic capacity with a general cohort of web users. This is the first study that employs such an algorithm as a diagnostic tool.

We use PD as a model to study the feasibility of remote diagnosis among web users independent of the content of typed text. Parkinson's disease is a progressive neurodegenerative disease diagnosed based on clinical findings (bradykinesia with resting tremor, muscle rigidity, or impaired gait). Remote diagnosis of PD individuals who have not actively sought medical care represents a new frontier with novel ethical challenges for medicine and for the collaborative potential of medicine and technology companies.

Methods

Interactions of web users with the Bing search engine (Microsoft Corporation, Seattle, USA) are routinely recorded to improve the performance of the engine, as described in the Terms of Service. We collected data between January 2016 and April 2017 from users in the United States querying the search engine. These data included an anonymized user identifier, contents of the search terms, locations of the mouse pointer over time, mouse clicks, and interactions with the keyboard (such as typing errors, time to first keyboard event) during the search process. This study was approved by the Behavioral Sciences Research Ethics Committee of the Technion (Haifa, Israel) approval number 2018-032. The methods were carried out in accordance with the relevant guidelines and regulations. Obtaining informed consent is not applicable to this study as data were collected using anonymized user identifiers. For ethical reasons, no attempt was made to identify users and researchers did not have access to any database linking anonymized user identifiers with IP addresses.

The relevant textual content of the queries was used to sort individuals into one of the two groups on which we trained the classifier (*PD* and *control groups*, see below). It was also used to create groups with expected higher rates of PD (*unclassified* and *at-risk* groups, see below). This content was not used in order to classify users as having PD in the three tested groups (*general population, unclassified,* and *at-risk* groups, see below). The classifier was not trained to distinguish PD from other types of parkinsonism (such as drug-induced).

Queries extraction process

Restricting our focus to queries containing a self-referral (i.e., "I have. . .," "My. . .") limited the number of examined users, but complied with our underlying assumption and should improve the reliability of our analysis.

For each of the study groups, we describe the specific adaptation employed to identify the relevant set of users:

1 *PD* group: Including only users for which at least one of the queries contained one of the following search phrases: "I have Parkinson's" (or other possible equivalent forms, e.g., "I've PD"); "I have been diagnosed with Parkinson's" (or other possible equivalent forms, e.g., "I've been diagnosed with PD"); "My Parkinson's" (or other possible equivalent forms, e.g., "My PD").
2 *Control* group**:** Including only users for which at least one of the queries contained one of the following: "My husband/wife/spouse has Parkinson's" (or other possible equivalent forms, e.g., "My husband has PD"); "Husband/wife/spouse with Parkinson's" (or other possible equivalent forms, e.g., "Husband with PD").
3 *Unclassified* group: To constrain the possibility of retrieving users searching for information regarding individuals with the last name "Parkinson," we considered only queries containing health-related information. We included only queries containing one of the following words: treatment, cure, disease, medication, medicine, drug, pain, doctor or diet. Additionally, we considered queries regarding Parkinson medicines that were FDA approved in the US, including Sinemet, Parcopa, Rytary, Duopa, Stavelo, Mirapex, Requip, Apokyn, Neupro, Eldepryl, Zelapar, Carbex, Azilect, Tasmar, Comtan, Symmetrel, Artane, and Cogentin. To discard queries by medical students regarding Parkinson's disease, we eliminated all queries containing the words patient, client or nurse.
4 *At-risk* group: To detect users searching with medical queries regarding a symptom that increases the likelihood of prodromal PD,[10] we searched for queries containing self-description ("I have," "I've," "I am," "I'm," "My"), and one of the following: brother or sister with Parkinson's, loss of smell, olfactory loss, constipation, somnolence, daytime sleepiness, erectile dysfunction, erectile, impotence, urinary dysfunction, depression, screaming during sleep, kicking during sleep, night terror, anosmia, hypotension, or low blood pressure.

Feature extraction

Mouse-tracking data consists of a list of time-stamped horizontal (x) and vertical (y) coordinates of the mouse cursor locations on the browser page during the interactions of a user.

We represented these data by extracting manually defined features (see Data S1). Keyboard attributes included the number of spelling mistakes in the query text and the time from when the results page was shown to the first click on the keyboard. To enrich this feature set, we further employed an automated supervised long short-term memory (LSTM, see below) model, which learns additional machine-generated features from the time-stamped raw data. This novel technique significantly improved the model performance and distinguishes our method from previous studies[7,11] in which only handcrafted features were used.

In order to automatically extract features from the raw mouse-tracking data, we predicted event-level mouse movements. Our automatic feature extraction task poses two challenges. First, in our setting, the mouse events are sampled in a nonuniform manner, and, therefore, the time gap between consecutive mouse events needs to be taken into account. Second, the sampling of the mouse events is

noisy and results in missing and corrupted data. To deal with the first problem, we predicted not only the next mouse position, but also the next event time stamp, allowing our algorithm to derive the mouse velocity, acceleration, and higher-level derivatives as intermediate quantities. To cope with the problem of noisy and corrupted data, we used robust scaling, that is, scaling the median of the mouse events to zero and the average absolute value of the 25th percentile and the 75th percentile to 1.

The scaled features were then inserted into a LSTM supervised model, followed by a fully connected layer, which attempted to predict the raw features of the next event. The model we employed for automatic feature extraction is the LSTM model of the keras library (https://keras.io/), using TensorFlow as the backend framework. We used the Adam optimizer and the mean squared error as the loss function. The dimension of the network was set to 1024, the number of training epochs was set to 7, the size of each batch to 256, and the learning rate was set to 0.002. We also examined the results while using other values of these parameters, and we report that the best results were achieved using the above parameter values for the hyperparameters. The outputs of the LSTM layer at the end of the session were chosen as the representative features of this session. Full details of the model and parameter choices for its training are given in https://dl.acm.org/citation.cfm?doxml:id=3269206.3269250.

Using both the manually and the automatically extracted features, we trained a classifier model that takes as an input a feature vector, corresponding to a user session, and produces an output of the likelihood that the user has PD (range 0–1). We examined multiple classification models (e.g., logistic regression, SVM) in the Microsoft TLC machine learning toolbox (https://azure.microsoft.com/en-us/services/machine-learning-studio/) and report the best results, achieved with the Boosted Decision Forest model, where the learning rate is set to 0.08 and the number of trees is 100. To provide a single score per user based on multiple sessions, we employ majority vote as an aggregation function to decide whether a user has the condition, based on the fraction of sessions with a score greater than 0.5.

Our code is publicly available on https://github.com/scientistl/Parkinson-unsupervised-features.

Statistical methods

Performance of classifiers was estimated using the receiver operating characteristic (ROC) and the area underneath the curve (AUC). ROC was chosen as the measure of performance for its suitability in the case of unbalanced datasets, such as the ones we analyzed. Differences between groups in the frequencies of users who were screened positive for PD were tested using the Fisher exact test. Means of slope of progression were compared using a two-tailed t-test and were corrected by the Bonferroni method. Values of P less than 0.05 were considered significant.

Results

Based on a content of a typed text, we first identified web users with PD and web users that served as controls without PD (Fig. 1, Step 1). The first group included users who disclosed having PD in their search query, by typing phrases such as "I have Parkinson's" (*PD group*, 281 individuals, 144,898 queries, 46–1988 queries per individual, median 612). The second group included web users who typed sentences such as "My husband has Parkinson's" (*control group*, 163 individuals, 182,594 queries, 56–2054 queries per individual, median 898). Web queries from two groups of users were used to train an automatic classifier. This classifier learned the complex behavioral features that can be used to identify individuals with PD, without relying on the content of the typed text.

The classifier was initially trained to calculate a score representing the likelihood that an individual query was produced by a user from the *PD group* (PD-likelihood score, range 0–1). To improve the performance of the classifier, we then grouped for each user the above calculated PD-likelihood scores of all her/ his single queries (46–2054 queries per user), and applied a majority vote aggregation criterion by calculating the fraction of sessions with a score greater than 0.5. If more than 50% of the queries received a PD-likelihood score greater than 0.5 this individual was classified as having PD. The resultant ROC reveals that the AUC for this model is 0.93, with true positive rates (sensitivity) of 0.93, 0.84, 0.62, and 0.45 corresponding to false positive rates (1-specificity) of 0.2, 0.1, 0.05, and 0.01.

The three most informative features of the prediction model (as reported by TLC) were the: (1) time elapsed from the moment results were presented to the user until first interaction event with the web page, (2) the average time elapsed between every two consecutive mouse positions (also referred to as average dwell time), and (3) one of the automatically learned features.

Following the training of the classifier, we applied this tool to a large cohort of web users, over the age of 40 years inferred from their date of birth as reported at registration to Bing (Fig. 1, Step 2). Users were included in this cohort only on the basis of their reported age (*general population group,* which comprised 1,490,987 individuals, 599,750,266 queries, number of queries per individual 42–4894, median 401). We used an estimated false positive rate of 0.01 paired with an estimated sensitivity of 0.45. These values yields an estimated positive predictive value

(PPV) of 25%, assuming a Parkinson's prevalence of 0.75% in this age group.[12–14] In this large cohort from the general population, 17,843 of 1,490,987 (1.2%) individuals were classified as having PD.

To validate the capability of the classifier to distinguish PD from non-PD, we then applied this tool to individuals from two additional groups of web users. We hypothesized that in these two groups an increased prevalence of affected individuals would be observed. Although the PPV of the classifier in these two groups is likely to be higher than in our *general population group*, due to higher *a priori* probability of PD,[10] we used the same threshold of classification (therefore, with an underestimation of the true prevalence of PD in these groups). Also here, the content of the typed text was only used to sort individuals to one of the studied group but was not used to classify individual as having PD or not.

The first validation group included users who repeatedly typed specific search terms related to PD (*unclassified group*, 804 individuals, 634,120 queries, number of queries per individual 49–2588, median 780). In this group, a significantly greater proportion of individuals were classified as users with PD, 518/804 (64.4%, Fisher exact *P*-value < 0.00001, relative to *general population group*). This result is consistent with previous work indicating that the majority of users who enter queries regarding medical conditions actually suffer from these conditions [15].

The second validation group included users over the age of 40 years who disclosed in their search query having a non-motor symptom of the disease and/or a sibling with PD (*at-risk group* 1,064 individuals, 713,802 queries, number of queries per individual 52–3018, median 664). Among this *at-risk group*, the model classified 57/1064 (5.3%, Fisher exact *P*-value < 0.00001, relative to *general population group*) individuals as ones with PD (see Data S2 for the number of individuals screened positive per each of the reported risk factor). In this group, 6/64 (10%, Fisher exact *P*-value < 0.0005, relative to *General population group*) of web users declared that one of their siblings had the disease were classified as having PD.

Finally, to further verify the validity of the classification (Fig. 1, Step 3), and to explore the possibility that this classifier can be used to follow patients longitudinally, we plotted over time, for each web user separately, 10% of her/his PD-likelihood scores that had the lowest values (one per query). We then fitted, to these 10% least 'parkinsonian' sessions of each user, a regression line in order to determine the slope of progression over time. In all of the three studied groups (*general population, unclassified and at-risk groups*), the mean of slope values of individuals who were classified as having PD were not found to be dissimilar to the mean of slope values of the *PD* group that was used as a reference (Fig. 2, two-tailed

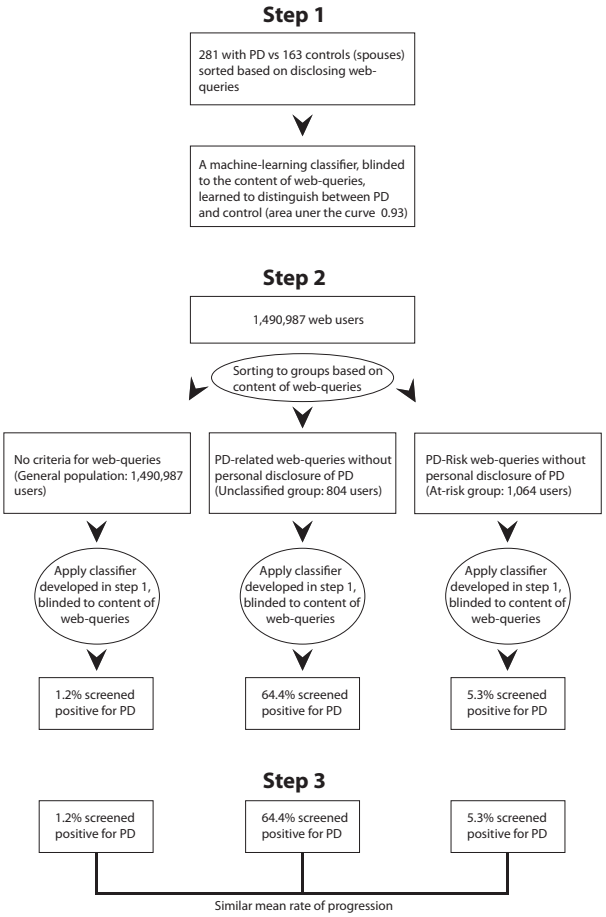

Figure 1. Flowchart and the main findings of the study

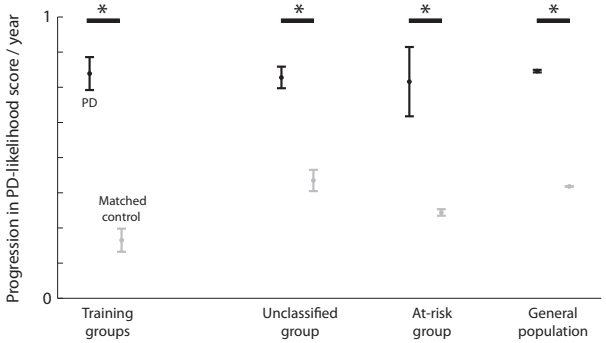

Figure 2. Progression of PD in individuals from the different study groups. Means and standard errors are shown for individuals screened positive (black) or negative (gray) for PD in each of the studied groups. Black horizontal lines with asterisks represent significant differences (*t*-test $P < 0.05$) between populations. The four populations screened positive for PD are statistically indistinguishable from each other, and the same applies for the four populations screened negative for PD. See Data S3 for exact *P*-values

t-test with Bonferroni correction, P-values ≥ 0.1). Similarly, the mean of slope values of individuals classified as not having PD from all study groups, were not dissimilar to the *control group* (two-tailed t-tests with Bonferroni correction, P- values ≥ 0.1). In contrast, for each of the studied groups the mean of progression slopes significantly differed between individuals who were classified as having PD and those from the same group that were classified as not having the disease (two-tailed t-test with Bonferroni correction, P-values < 0.05). These results indicate that our classification results were nonrandom and reflect the actual diagnostic medical condition of individuals. These results also demonstrate the potential of the classifier to document disease progression.

Discussion

Our study validates the successful application of a machine learning algorithm to large-scale, web-based screening for PD. As a screening tool, the algorithm has an acceptable PPV (estimated as 25% here). This PPV is not inferior to that of other commonly used screening tests.[16] A clinician examining a patient would diagnose PD more accurately once motor signs become apparent.[17] It has been shown that a simple solicited finger-tapping task can predict increased risk PD[18] but it remains to be determined in future studies if an automatic algorithm can perform better than a human clinician in identifying prodromal PD when motor signs are subtle or absent. In this study we did not attempt to distinguish prodromal from motor PD.

To a reasonable degree, a human PD expert is capable of distinguishing PD from other neurodegenerative diseases (e.g., Multiple System Atrophy).[19] Diagnostic accuracy drops from 83.9% to 73.8% when nonexperts perform the diagnosis.[20] In our study, we could not collect data regarding the diagnostic process of individuals in the PD group that served to train the classifier. It is reasonable to assume that some users in this group had other Parkinsonian syndromes. In order to establish, in future studies, the true PPV and negative predictive value (NPV) of a web screening tool, a validation of reported diagnosis using other medical databases (or physical examination of users) would be required. For ethical reasons, we could not do it in this fully anonymized study design. Due to all these limitations, in this study we did not train our classifier to distinguish between PD and other neurodegenerative diseases. We suspect that such a task will be feasible as well in future algorithm refining and development.

The ability to identify patients with PD, and to document their motor decline, based on interactions with digital devices such as computer keyboards[8,11,21] or mobile touch screens[22,23] was demonstrated in previous studies. These algorithms, however, have never been validated in real-life. Unlike in these aforementioned studies, the present study provides validation of its algorithm tested in a large cohort of web users considered to be at low or high risk of actual PD. Increased frequency of PD cases detected by our algorithm among web users who searched for PD-related terms, reported a prodromal symptom of the disease or a blood-related relative with the disease enhanced the validity of this method. This conclusion is also supported by the longitudinal progression of PD-related features, as defined by the algorithm, which was much faster among these with suspected PD in all tested groups.

Currently, an early diagnosis of PD would not change the progression of the disease but may actually benefit an individual with unexplained impairment of quality of life due to constipation, insomnia, depression, or other nonspecific symptoms which would not have otherwise directed him to appropriate medical/neurological attention. In addition, knowledge regarding a diagnosis would lead some individuals to seek solutions (e.g., participation in clinical trials testing disease-modifying therapies), and may shape personal decisions regarding the future. Parkinson's disease web diagnosis could add to the knowledge of the disease, for example, through future studies comparing disease prevalence in different populations, or examining environmental factors that reduce or increase the risk for PD or that contribute to disease progression. These potential benefits of unsolicited diagnosis, such as the one demonstrated here, should be considered in light of the ethical issues.

Unsolicited diagnosis raises important concerns. There has been some discussion of the ethical issues raised by unsolicited diagnosis in medical fields such as neurology,[24] dermatology,[25] and psychiatry,[26] in cases where diagnosis is made by "bystander" physicians uninvolved in the "patient's" usual care. Issues that have been addressed include the risk of misdiagnosis, stigma to patients, confidentiality, overdiagnosis, and overtreatment.[26,27] These issues are all relevant for the future drafting of ethical guidelines for unsolicited web detection of diseases. Studies of public attitude to the disclosure of unsolicited diagnoses indicate that it is generally favorable,[28] even in the case of serious disease that cannot be modified, or in regards to information that is of uncertain clinical significance.[29]

Our findings highlight the urgency in the need to establish ethical guidelines for technology companies and researchers involved in unsolicited web-derived diagnoses. For obvious ethical reasons, we did not attempt identifying subjects in this study. Our objective focused on demonstrating that with the accelerating development of remote, unsolicited web-based diagnosis ethical dilemmas move outside of the sole area of responsibility of the

medical profession to encompass technology companies that develop capabilities to collect and analyze user information on a massive scale. The absence of an ethical framework dealing with this pertinent issue could harm both users and commercial companies, and has far-reaching implications for the current practice of medicine. Collaboration between the medical community, the public and the leading technology companies is required to develop an ethical framework and guidelines for the use of web-based diagnostic tools, and for informing users of their results. Such collaboration could improve users' well-being while maintaining their rights to privacy, their ability to receive clinically useful information, their autonomy to choose between different possible courses of action, and most notably, their right not to know.

Acknowledgments

The authors thank Prof. Tamir Ben-Hur for his comments that helped improving this manuscript.

Author Contributions

All authors made substantial contributions to this work and approved the submitted version. BY, LA, and EYT contributed to conception and design of the study, acquisition, analysis, and interpretation of data, and creation of new software used in the work. BY and EYT also drafted the work. OP contributed to analysis, drafted this work, and substantially revised it. DA contributed to conception and design of this work, statistical analysis, drafted this work and substantially revised it.

Conflicts of Interest

Brit Youngmann, Liron Allerhand, and Elad Yom-Tov are employees of Microsoft, owner of the Bing search engine, and their part of the work described in this paper was partially performed as part of the authors' salaried employment. Ora Paltiel and David Arkadir report no conflicts of interest.

Data Availability Statement

The datasets generated during the current study are available from the Dr. Elad Yom-Tov email: eladyt@microsoft.com on reasonable request.

References

1. Ginsberg J, Mohebbi MH, Patel RS, et al. Detecting influenza epidemics using search engine query data. Nature 2009;457:1012–1014.
2. Althouse BM, Ng YY, Cummings DAT. Prediction of dengue incidence using search query surveillance. PLoS Negl Trop Dis 2011;5:e1258.
3. Shaman J, Karspeck A, Yang W, et al. Real-time influenza forecasts during the 2012–2013 season. Nat Commun 2013;4:2837.
4. Ayers JW, Ribisl KM, Brownstein JS. Tracking the rise in popularity of electronic nicotine delivery systems (electronic cigarettes) using search query surveillance. Am J Prev Med 2011;40:448–453.
5. White RW, Tatonetti NP, Shah NH, et al. Web-scale pharmacovigilance: listening to signals from the crowd. J Am Med Inform Assoc JAMIA 2013;20:404–408.
6. Paparrizos J, White RW, Horvitz E. Screening for pancreatic adenocarcinoma using signals from web search logs: feasibility study and results. J Oncol Pract 2016;12:737–744.
7. White RW, Horvitz E. Evaluation of the feasibility of screening patients for early signs of lung carcinoma in web search logs. JAMA Oncol 2017;3:398–401.
8. White RW, Doraiswamy PM, Horvitz E. Detecting neurodegenerative disorders from web search signals. Npj Digit Med 2018;1:8.
9. Allerhand L, Youngmann B, Yom-Tov E, Arkadir D. Detecting Parkinson's Disease from interactions with a search engine: Is expert knowledge sufficient? [Internet]. ArXiv180501138 Cs 2018;[cited 2018 Aug 31] Available from: http://arxiv.org/abs/1805.01138
10. Berg D, Postuma RB, Adler CH, et al. MDS research criteria for prodromal Parkinson's disease: MDS Criteria for Prodromal PD. Mov Disord 2015;30:1600–1611.
11. Giancardo L, Sánchez-Ferro A, Arroyo-Gallego T, et al. Computer keyboard interaction as an indicator of early Parkinson's disease. Sci Rep 2016;6:34468.
12. de Rijk MC, Tzourio C, Breteler MM, et al. Prevalence of parkinsonism and Parkinson's disease in Europe: the EUROPARKINSON Collaborative Study. European Community Concerted Action on the Epidemiology of Parkinson's disease. J Neurol Neurosurg Psychiatry 1997;62:10–15.
13. de Rijk MC, Breteler MM, Graveland GA, et al. Prevalence of Parkinson's disease in the elderly: the Rotterdam Study. Neurology 1995;45:2143–2146.
14. de Lau LM, Breteler MM. Epidemiology of Parkinson's disease. Lancet Neurol 2006;5:525–535.
15. Yom-Tov E, Borsa D, Hayward AC, et al. Automatic identification of web-based risk markers for health events [Internet]. J Med Internet Res 2015;17[cited 2018 Jul 9].
16. Kerlikowske K, Grady D, Barclay J, et al. Positive predictive value of screening mammography by age and family history of breast cancer. JAMA 1993;270:2444–2450.
17. Hughes AJ, Daniel SE, Kilford L, Lees AJ. Accuracy of clinical diagnosis of idiopathic Parkinson's disease: a

clinico-pathological study of 100 cases. J Neurol Neurosurg Psychiatry 1992;55:181–184.
18. Noyce AJ, R'Bibo L, Peress L, et al. PREDICT-PD: An online approach to prospectively identify risk indicators of Parkinson's disease. Mov Disord Off J Mov Disord Soc 2017;32:219–226.
19. Litvan I, Goetz CG, Jankovic J, et al. What is the accuracy of the clinical diagnosis of multiple system atrophy?: a clinicopathologic study. Arch Neurol 1997;54:937–944.
20. Rizzo G, Copetti M, Arcuti S, et al. Accuracy of clinical diagnosis of Parkinson disease: a systematic review and meta-analysis. Neurology 2016;86:566–576.
21. Adams WR. High-accuracy detection of early Parkinson's Disease using multiple characteristics of finger movement while typing. PLoS ONE 2017;12:e0188226.
22. Arroyo-Gallego T, Ledesma-Carbayo MJ, Sanchez-Ferro A, et al. Detection of motor impairment in parkinson's disease via mobile touchscreen typing. IEEE Trans Biomed Eng 2017;64:1994–2002.
23. Iakovakis D, Hadjidimitriou S, Charisis V, et al. Touchscreen typing-pattern analysis for detecting fine motor skills decline in early-stage Parkinson's disease. Sci Rep 2018;8:7663.
24. Lim ECH, Quek AML, Seet RCS. Putting ethics on the spot in neurology. Nat Clin Pract Neurol 2009;5:136–137.
25. Preller G, Salloch S. Melanoma in the shopping mall: A utilitarian argument for offering unsolicited medical opinions in informal settings. Bioethics 2018;32:193–198.
26. Mitchell EW. The ethics of unsolicited diagnosis of mental disorder in acquaintances: benefits and dangers. Psychiatrist 2011;35:297–301.
27. Volk ML, Ubel PA. Better off not knowing: improving clinical care by limiting physician access to unsolicited diagnostic information. Arch Intern Med 2011;171:487–488.
28. Zwitter M, Nilstun T, Knudsen LE, et al. Professional and public attitudes towards unsolicited medical intervention. BMJ 1999;318:251–253.
29. Middleton A, Morley KI, Bragin E, et al. Attitudes of nearly 7000 health professionals, genomic researchers and publics toward the return of incidental results from sequencing research. Eur J Hum Genet EJHG 2016;24:21–29.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Data S1. Supplemental methods: Manually defined features used by the prediction.
Data S2. Supplemental results: Number of individuals screened positive per each of the reported risk factor.
Data S3. Supplemental results: *P*-values for differences between progression slopes.

