

Assessment of Parkinson’s Disease Medication State through Automatic Speech Analysis

Anna Pompili¹, Rubén Solera-Ureña¹, Alberto Abad^{1,2}, Rita Cardoso^{3,5,6}, Isabel Guimarães^{5,6,7}, Margherita Fabbri⁸, Isabel P. Martins^{4,5}, Joaquim Ferreira^{3,5,6}

¹INESC-ID, Lisbon, Portugal

²Instituto Superior Técnico, Universidade de Lisboa (UL), Portugal

³Laboratory of Clinical Pharmacology and Therapeutics, Faculdade de Medicina, UL, Portugal

⁴Laboratório de Estudos de Linguagem, Faculdade de Medicina, UL, Portugal

⁵Instituto de Medicina Molecular, Lisbon, Portugal

⁶CNS - Campus Neurológico Sénior, Torres Vedras, Portugal

⁷Alcoitão School of Health Sciences, Santa Casa da Misericórdia de Lisboa, Portugal

⁸Departments of Clinical Pharmacology and Neurosciences, CIC 1436, NS-Park/FCRIN network and NeuroToul COEN Center, INSERM, CHU of Toulouse and University of Toulouse, France

anna@hlt.inesc-id.pt, rsolera@hlt.inesc-id.pt, alberto.abad@inesc-id.pt

Abstract

Parkinson’s disease (PD) is a progressive degenerative disorder of the central nervous system characterized by motor and non-motor symptoms. As the disease progresses, patients alternate periods in which motor symptoms are mitigated due to medication intake (ON state) and periods with motor complications (OFF state). The time that patients spend in the OFF condition is currently the main parameter employed to assess pharmacological interventions and to evaluate the efficacy of different active principles. In this work, we present a system that combines automatic speech processing and deep learning techniques to classify the medication state of PD patients by leveraging personal speech-based bio-markers. We devise a speaker-dependent approach and investigate the relevance of different acoustic-prosodic feature sets. Results show an accuracy of 90.54% in a test task with mixed speech and an accuracy of 95.27% in a semi-spontaneous speech task. Overall, the experimental assessment shows the potentials of this approach towards the development of reliable, remote daily monitoring and scheduling of medication intake of PD patients.

Index Terms: Parkinson’s Disease speech, ON-OFF medication state, automatic assessment

1. Introduction

Parkinson’s disease (PD) is a progressive degenerative disorder of the central nervous system characterized by motor and non-motor symptoms. The cardinal motor signs of PD include the characteristic clinical picture of resting tremor, rigidity, bradykinesia, and impairment of postural reflexes, while non-motor symptoms include cognitive disorders, and sleep and sensory abnormalities. Motor symptoms of PD influence also the speech production of language. Dysarthria, which is characterized by a weakness, paralysis, or lack of coordination in the motor-speech system, is typically observed in PD patients and affects respiration, phonation, articulation and prosody. As a consequence, the main deficits of PD speech are loss of intensity, monotony of pitch and loudness, reduced stress, inappro-

priate silences, short rushes of speech, variable rate, imprecise consonant articulation and harsh and breathy voice. Both motor symptoms and speech impairments slowly worsen during the disease with a nonlinear progression. Patients in advanced stages typically present more severe speech abnormalities, with voice disorders and articulatory deficits being the most prevalent symptoms. At the final stage of the disease, articulation is frequently the most impaired feature [1–3]. As the disease progresses, patient alternate periods in which motor symptoms are mitigated due to medication intake (ON state), and periods with motor complications (OFF state). The time that patients spend in the OFF condition is currently the main parameter used to assess pharmacological interventions.

In this scenario, it is of paramount importance to investigate whether speech characteristics can be used as a bio-marker to monitor and track ON-OFF conditions of PD patients. If this would be possible, then speech cues may bring further insights into the dynamics of the disease over time, not only by providing accurate information on medication states and alterations, but also by investigating its correlation with motor fluctuations. However, while the effects of dysarthria in PD have been deeply investigated, both in drug-naïve state and under pharmacological treatment [4–9], it is still not clear whether speech analysis could be helpful to differentiate between ON-OFF states induced by pharmacological intake. In fact, the clinical literature presents controversial results on the effect of Levodopa on speech abilities [10–15], while no computational studies addressing this topic have been identified in the literature.

To the best of our knowledge, this is the first work that investigates the problem of automatically identifying medication states of PD patients from speech. We propose a system that combines speech processing and deep learning techniques to perform automatic classification of the medication state by leveraging personal speech-based bio-markers. We devise a speaker-dependent experimental setup that uses the FraLusoPark corpus [16]. This dataset contains recordings from 74 native Portuguese PD patients performing several vocal tasks, both in drug-naïve state and under pharmacological treatment. The results show that it is possible to classify ON-OFF condition with an accuracy of 90.54% using speaker-dependent

This work has been supported by national funds through Fundação para a Ciência e a Tecnologia (FCT) with reference UIDB/50021/2020.

models. Moreover, it was found in previous work that the type of vocal task has a relevant impact on the capability of automatic approaches to characterize speech impairments in PD patients [6, 17, 18]. When we extend such a per-speech task analysis to this work, we find that the accuracy of the system improves to 95.27% on a semi-spontaneous storytelling task. These results are of extreme relevance, not only because they represent a first contribution to the automatic assessment of the medication state of PD patients. They also confirm that it is possible to obtain reliable models using a relatively reduced amount of patients’ speech data. Altogether, we consider that this is a relevant achievement towards the development of reliable, personalized and non-invasive health solutions for the daily monitoring and scheduling of medication intake of PD patients. This technology could be deployed at patients’ homes and remotely provide the clinicians with constant, updated information.

The rest of this work is structured as follows: Section 2 introduces the relevant state of the art. In Section 3 and 4, we present the dataset used in this study and a description of our methodology. Experimental results are presented in Section 5 and the paper ends with conclusions in Section 6.

2. Related work

The effects of Levodopa treatment on speech abilities have been analyzed considering different dimensions of voice quality. Several studies focused their attention on phonatory and articulatory abilities. Goberman *et al.* [19] found small differences between groups of patients before and after medication, which were pointing to phonatory improvements. Jacobi *et al.* [20] analyzed the Vowel Articulation Index (VAI) in 10 participants of different languages, finding that VAI was not significantly affected by pharmacological treatment. On the other hand, Okada *et al.* [10] reported significantly expanded Vowel Space Area (VSA) after Levodopa treatment of 21 PD patients, contrarily to the findings by Poluha *et al.* [11] on a corpus of 10 subjects. Some authors shifted their attention towards the analysis of prosodic deviations. Using a reading task composed of four complex sentences, Skodda *et al.* [12] analyzed intonation and speech rate in a group of 138 patients. The authors found significant changes from the first to the last sentence of the passage, but generally, no effect of Levodopa administration was found on the reading of the entire task. In this respect, De Letter *et al.* [13] found no differences in the speech rate of 25 patients reading a text passage in OFF and ON states. Finally, when investigating the effects of Levodopa on developmental stuttering, Anderson *et al.* [14] reported an increase of disfluencies in the spontaneous speech of 1 patient during the ON period. This is in disagreement with the findings of Im *et al.* [15], which examined 51 patients while reading a standard passage. In fact, the authors found an improvement of disfluency in the ON state that was proportional to the severity of speech impairments presented in the OFF state. The discrepancies found in the studies described above may depend on many factors, like different methodological approaches or the small number of samples observed in some studies. Overall, it is rather consensual that Dysarthria is considered to be responsive to treatments. Nevertheless, there exists also the hypothesis that pharmacological intake may induce additional speech deficits caused by side effects of the treatment itself [21].

While speech and language technologies may potentially provide valuable contributions to the analysis of patients’ states, to the best of our knowledge, this is the first study that investigates a computational approach to this problem. In this re-

Table 1: *Duration by task and medication state of the FraLusoPark dataset (hh:mm:ss) after automatic pre-processing.*

Task	ON	OFF
/a/	00:31:22	00:35:22
MPT	00:39:22	00:38:24
DDK	00:32:36	00:30:50
Reading 10 words	00:10:34	00:09:52
Reading 10 sentences	00:18:24	00:17:56
Reading text	00:40:04	01:27:56
Reading prosodic sentences	01:22:21	00:41:08
Story telling	00:37:59	00:41:14
Conversation	02:40:38	02:51:21
Total	07:18:17	07:36:19

spect, our work brings the following contributions. First, it is a fully-automatic method purely based on acoustic bio-markers of speech. Second, we explore the use of deep architectures in a challenging domain like the health area, characterized by relatively constrained datasets with respect to the ones typically used with these models. Third, by exploiting a relatively large corpus composed of 74 PD patients performing several tasks, we show that is possible to achieve reliable personalised models to identify ON-OFF medication states, particularly in semi-spontaneous vocal production tasks.

3. Corpus description

The corpus used in this study is the FraLusoPark database [16], which contains the recordings of 139 European Portuguese speakers. The control group, composed of 65 healthy volunteers, is age-matched and sex-matched with the PD group, composed of 74 subjects. The patient group contains 38 male with an average age of 65 years (± 11.9), and 37 female, with an average age of 70 (± 8.5) years. Each patient was recorded twice on the same day, OFF medication (i.e.: at least 12 hours after withdrawal of all anti-Parkinsonian drugs), and ON medication (i.e.: following at least 1 hour after the administration of the usual medication). All the recordings were performed under exactly the same environmental conditions.

Participants were required to perform 9 speech production tasks with an increasing complexity in a fixed order: (1) three repetitions of the sustained phonation of the vowel /a/, (2) two repetitions of the maximum phonation time (vowel /a/ sustained as long as possible), (3) oral diadochokinesia (repetition of the pseudo-word *pataka* at a fast rate for 30 s.), (4-5) reading aloud of 10 words and 10 sentences, (6) reading aloud of a short text (“The North Wind and the Sun”), (7) reading aloud of a set of sentences with specific prosodic properties, (8) storytelling speech guided by visual stimuli, and (9) engaging in a spontaneous conversation for around three minutes. Table 1 shows the duration of the corpus by task and medication state after removal of pauses and therapists’ interventions.

4. Methodology

Our approach to the automatic classification of ON-OFF states relies on three main stages. The first one deals with the automatic removal of non-speech segments and therapists’ interventions. Then, acoustic-prosodic markers are computed from the clean speech signals. Finally, PD patients’ medication state is modeled with deep neural networks.

4.1. Pre-processing and segmentation

We describe here our automatic approach to remove speech segments corresponding to silent pauses and instructions of the speech language therapists. Each audio file in this corpus corresponds to the continuous, unsegmented recording of a task. In our previous work with this dataset [17], the recordings of the control group were processed semi-automatically to generate segmentation annotations with information of speech/non-speech segments and speaker (patient/clinician) turn. These data are exploited to generate the same kind of annotations for the Patient recordings, using a fully-automatic approach based on speech/non-speech segmentation (SNS) followed by speaker identification. The SNS segmentation is based on a simple finite state machine encoding heuristic rules about the minimal segment duration and state-changing thresholds, and takes as input the likelihoods provided by a bi-Gaussian model of the log-energy distribution computed for each recording. In the speaker identification step, we follow the classical GMM-MAP modeling approach [22], in which a universal background model (UBM) is first trained and then adapted to a target speaker using some enrolment data. In this case, we train the UBM using all the speech segments in the control group recordings (07h 19m 27s) and a therapist model using the therapist speech segments in the control group recordings (07m and 40s). Finally, for each speech segment in the Patient recordings, the therapist vs. UBM likelihood ratio is computed and compared to a decision threshold Θ . In this work, we are specially concerned with the removal of the therapist’s speech segments, while, up to some extent, we can afford to eliminate some data of patient’s speech. For that reason, the decision threshold was empirically set to minimize the therapist insertion errors according to experiments conducted in the Control group recordings.

4.2. Feature extraction

After the removal of pauses and therapists’ interventions, all the clean speech recordings are parameterized in a frame basis. Three different feature sets were used in this work. The first two sets, henceforth referred to as MFCC and MFCC+ Δ s, consist of 13 Mel-Frequency Cepstral coefficients (MFCC) in the former case, augmented with the corresponding 13 delta coefficients in the latter. These features were extracted with HTK [23] using a window size of 25 ms and a window period of 10 ms. The third feature set, referred to as eGeMAPS, consists of the 26 low-level descriptors (LLD) of the eGeMAPS feature set [24]. This is a compact set of acoustic-prosodic features, well-known for their usefulness in a wide range of paralinguistic tasks. These features were extracted with the openSMILE toolbox [25] using a frame period of 10 ms and window sizes of 20 and 60 ms depending on the specific LLD [24]. The feature vectors are finally normalized (zero mean and unit variance) on a per-speaker basis.

4.3. Medication state modeling

The task of assessing PD patients’ medication state (ON-OFF) from their speech can be stated as a binary classification problem. In a preliminary phase, we experimented a simple model like GMMs, which provided poor results. Thus, we use deep neural networks (DNN) for assigning every input frame in the patient’s utterance with a probability of that speech sample corresponding to either states ON or OFF. A final decision for the whole utterance is obtained as the mean value of the individual probabilities. One fundamental aspect in the definition of the

approach followed to model the medication state regards the dependency of this task on the speakers’ personal characteristics. This comes down to determining whether there exists a set of acoustic-prosodic cues shared by speakers with different gender, ages, and health conditions (including PD development stage) that can be leveraged to develop robust models for this task. Thus, we conducted a preliminary series of experiments in a speaker-independent setup. Several machine learning techniques (DNN and SVM) and feature types (MFCC, eGeMAPS, i-vectors and X-vectors) were used to train speaker-independent models that achieved limited performance results up to 65% utterance-level accuracy. These results, along with the nature of the task itself, justify the adoption of the speaker-dependent approach that is presented in detail below.

4.3.1. Data partitioning

The pre-processed FraLusoPark dataset is divided into three different subsets for training, development (validation) and test purposes. In a speaker-dependent approach, all these three subsets contain recordings from all the 74 PD patients. Namely, we divide the pre-processed dataset by speech tasks, as following:

- Train: MPT, DDK, Reading 10 words, Reading prosodic sentences, and Conversation.
- Development: Reading 10 sentences.
- Test: /a/, Reading text, and Story telling.

The training set contains 740 files (10 files per patient, 5 ON and 5 OFF) and 3983327 samples, which amounts to 55.6% of the whole pre-processed dataset in terms of utterances. The development set contains 148 files (2 per patient) and 217956 samples (11.1% of the dataset). The test set contains 444 files (6 per patient) and 1166119 samples, (remaining 33.3% of the dataset). On average, there are approximately 53800, 2900 and 15800 samples per speaker in these 3 subsets (around 9, 0.5 and 2.6 minutes of speech), respectively. Overall, the number of samples corresponding to the ON and OFF states is well-balanced (2629635 vs. 2737767), although we could observe some notable differences for some speakers, specially in the training subset.

4.3.2. Medication state classifiers

Different architectures of feedforward networks were assessed for the goal of learning the relationships between input feature vectors and ON-OFF medication states. Specifically, architectures with 1 to 3 hidden layers with 32 to 512 nodes were validated. ReLU and sigmoid activation functions are used in the hidden and output layers, respectively. The learning rate α is validated, whereas a moderate fixed batch size (32) is used to limit the search space in the validation process. Also, different input context sizes were assessed in our experiments (1, 5, 11, and 15 frames), together with the application of Principal Component Analysis (PCA) for dimensionality reduction. The selection of the optimal values for the free (hyper)parameters was performed based on validation results using the development subset. The final results are computed on the test subset. All the results in this work are reported in terms of accuracy (Acc) at a file/utterance-level.

In our speaker-dependent approach, the final system consists of 74 individual models, one for each speaker. However, they all share the same configuration (architecture and training hyperparameters). The reason behind this is that just 2 speech recordings from each patient are available for validation purposes. Thus, validation on a per-speaker basis is not feasible.

Table 2: Optimal model configurations and results (utterance-level Acc - %) for speaker-dependent medication state assessment.

Feature set	#Coefficients	Context	Input dim.	Architecture	α	Acc devel.	Acc test
MFCC	13	15	195	256, 128, 32, 1	0.01	93.92	88.74
MFCC+ Δ s	26	11	286	256, 128, 1	0.01	93.24	89.86
eGeMAPS	23	15	345	512, 128, 1	0.001	95.95	90.54
MFCC+PCA	13	15	95	512, 128, 1	0.03	91.89	90.09
MFCC+ Δ s+PCA	26	11	85	128, 64, 1	0.03	91.89	89.41
eGeMAPS+PCA	23	15	70	128, 64, 1	0.03	96.62	87.84

Instead, the average accuracy on the validation subset across all the patients is used to select the optimal global configuration that is then used to train 74 individual DNNs, one for each patient. The DNN classifiers were trained on several GPUs using TensorFlow [26].

Table 2 shows the optimal configuration for each of the six parameterizations used in this work: the three sets described in Section 4.2 with and without PCA. The number of coefficients per speech frame are 13, 26 and 23 for MFCC, MFCC+ Δ s and eGeMAPS, respectively. A context length of 15 frames was found optimal for the MFCC- and eGeMAPS-based systems, whereas the MFCC+ Δ s(+PCA) systems use a context of 11 frames. For the PCA-based systems, the input dimension corresponds to the number of principal components kept, which is deterministically selected as the n largest eigenvalues that explain a 95% of the total variance in the original input space.

5. Experimental results and analysis

Table 2 shows the validation and test results (utterance-level accuracies) of the DNN-based systems described above. These results reveal that the medication condition (ON-OFF) of the speakers can be inferred from the acoustic-prosodic information in their speech by means of automatic speech processing and machine learning procedures. The MFCC-based systems achieve accuracies of 88.74% and 90.09% in the test set, while the MFCC+ Δ s-based systems achieve comparable results of 89.86% and 89.41%. One reason to explain the competitive results achieved by the MFCCs alone may be the use of a larger input context size (15 frames), which compensates for the dynamic information provided by the delta MFCC features. The eGeMAPS-based systems achieve accuracies of 90.54% and 87.84%. The use of PCA leads to better results for the MFCC-based system, even with a substantial reduction in the dimensionality (51.3%). For the MFCC+ Δ s and the eGeMAPS-based systems the results are lower, however. We hypothesize that this may be due to the extreme dimensionality reduction in these two cases (70.3% and 79.7%, respectively). Finally, it is worth mentioning that the medication state is not equally assessed for all the patients. Since all of them contribute the same number of test utterances, the results in Table 2 coincide with the average per-speaker accuracies. The speaker variability, expressed in terms of the standard deviation of the per-speaker accuracies, ranges from 12.51 to 14.67.

The markedly different nature of the various speech tasks in the test subset justifies a separate analysis of the results in a per-task basis. Table 3 shows how the performances of the six models depend on the speech task: phonation of the vowel /a/, reading of a short text, and guided storytelling. The best results are achieved on the semi-spontaneous visually-guided storytelling task (accuracies between 92.57 and 95.27%), followed by the reading of short text task (accuracies between 86.49 and 91.22%) and the /a/ task (accuracies around 83.78% to 89.19%).

It is specially remarkable the performance of the eGeMAPS on the story telling task, with a 95.27% accuracy. As a conclusion, the results shown in Table 3 point to the benefit of using natural, (semi)spontaneous speech recordings for the automatic assessment of PD patients medication state. These observations are consistent with previous work on the relevance of speech tasks for automatic detection of Parkinson’s disease [17]. Overall, the results reported in this work have the potential to contribute to opening new avenues for the automatic distant monitoring of PD patients.

Table 3: Test results (utterance-level Acc - %) for speaker-dependent medication state assessment by task.

Feature set	/a/	Reading text	Story telling
MFCC	86.49	87.16	92.57
MFCC+ Δ s	85.14	91.22	93.24
eGeMAPS	89.19	87.16	95.27
MFCC+PCA	86.49	89.19	94.59
MFCC+ Δ s+PCA	85.14	89.19	93.92
eGeMAPS+PCA	83.78	86.49	93.24

6. Conclusions and future work

This work presents an approach that combines speech processing and deep learning techniques to perform automatic classification of the medication state of PD patients by leveraging personal speech-based bio-markers. We devise a speaker-dependent approach and investigate the relevance of different acoustic-prosodic feature sets. Test results show an accuracy of 90.54% in a mixed-speech task and a remarkable accuracy of 95.27% in a semi-spontaneous speech task. These results show the potentials of this approach towards the development of personalized and reliable systems for a daily, remote monitoring and scheduling of medication intake of PD patients.

Further work will pursue the acquisition of more speech data to perform more exhaustive and relevant experimentation, together with an analysis of how the proposed approach scales to new and larger corpora. Presumably, this would also open the door to new and more elaborated approaches using convolutional and/or recurrent neural networks, and to speaker-independent approaches. We plan to continue investigating that approach in combination with the use of speech embeddings (e.g., i-vectors and X-vectors) that may convey specific information relevant to the medication state classification task. Speaker-independent models may also serve as a good starting point for the development of more powerful speaker-dependent models by means of retraining them using speaker-specific speech data.

7. References

- [1] A. K. Ho, R. Ianse, C. Marigliani, J. L. Bradshaw, and S. Gates, "Speech impairment in a large sample of patients with Parkinson's disease," *Behavioural Neurology*, vol. 11, no. 3, pp. 131–137, 1999.
- [2] S. Sapir, A. Pawlas, L. Ramig, S. Countryman, C. O'Brien, M. Hoehn, and L. Thompson, "Speech and voice abnormalities in Parkinson disease: relation to severity of motor impairment, duration of disease, medication, depression, gender and age," *NCVS Status and Progress Report*, vol. 14, pp. 149–161, 1999.
- [3] J. A. Logemann, H. B. Fisher, B. Boshes, and E. R. Blonsky, "Frequency and Cooccurrence of Vocal Tract Dysfunctions in the Speech of a Large Sample of Parkinson Patients," *Journal of Speech and Hearing Disorders*, vol. 43, no. 1, pp. 47–57, 1978.
- [4] S. Skodda, "Aspects of speech rate and regularity in Parkinson's disease," *Journal of the Neurological Sciences*, vol. 310, no. 1–2, pp. 231–236, 2011.
- [5] T. Bocklet, E. Nöth, G. Stemmer, H. Ruzickova, and J. Rusz, "Detection of persons with Parkinson's disease by acoustic, vocal, and prosodic analysis," in *2011 IEEE Workshop on Automatic Speech Recognition & Understanding*, 2011, pp. 478–483.
- [6] J. Rusz, R. Cmejla, H. Ruzickova, and E. Ruzicka, "Quantitative acoustic measurements for characterization of speech and voice disorders in early untreated Parkinson's disease," *The Journal of the Acoustical Society of America*, vol. 129, no. 1, pp. 350–367, 2011.
- [7] J. A. Whitfield, C. Delong, A. M. Goberman, and M. Blomgren, "Fluency adaptation in speakers with Parkinson disease: a motor learning perspective," *International Journal of Speech-Language Pathology*, vol. 20, no. 7, pp. 699–707, 2018.
- [8] S. Skodda, W. Visser, and U. Schlegel, "Vowel articulation in Parkinson's disease," *Journal of Voice*, vol. 25, no. 4, pp. 467–472, 2011.
- [9] V. L. Hammen and K. M. Yorkston, "Speech and pause characteristics following speech rate reduction in hypokinetic dysarthria," *Journal of Communication Disorders*, vol. 29, no. 6, pp. 429–445, 1996.
- [10] Y. Okada, M. Murata, and T. Toda, "Effects of levodopa on vowel articulation in patients with Parkinson's disease," *Kobe Journal of Medical Sciences*, vol. 61, no. 5, pp. E144–54, 2016.
- [11] P. Poluha, H.-L. Teulings, and R. Brookshire, "Handwriting and speech changes across the levodopa cycle in Parkinson's disease," *Acta Psychologica*, vol. 100, no. 1–2, pp. 71–84, 1998.
- [12] S. Skodda, W. Grönheit, and U. Schlegel, "Intonation and speech rate in parkinson's disease: General and dynamic aspects and responsiveness to levodopa admission," *Journal of Voice*, vol. 25, no. 4, pp. e199–e205, 2011.
- [13] M. De Letter, P. Santens, M. De Bodt, P. Boon, and J. Van Borsel, "Levodopa-induced alterations in speech rate in advanced Parkinson's disease," *Acta Neurologica Belgica*, vol. 106, no. 1, pp. 19–22, 2006.
- [14] J. M. Anderson, J. D. Hughes, L. J. G. Rothi, G. P. Crucian, and K. Heilman, "Developmental stuttering and Parkinson's disease: the effects of levodopa treatment," *Journal of Neurology, Neurosurgery & Psychiatry*, vol. 66, no. 6, pp. 776–778, 1999.
- [15] H. Im, S. Adams, A. Abeyesekera, M. Pieterman, G. Gilmore, and M. Jog, "Effect of Levodopa on Speech Dysfluency in Parkinson's Disease," *Movement Disorders Clinical Practice*, vol. 6, no. 2, pp. 150–154, 2019.
- [16] S. Pinto, R. Cardoso, J. Sadat, I. Guimarães, C. Mercier, H. Santos, C. Atkinson-Clement, J. Carvalho, P. Welby, P. Oliveira, M. D'Imperio, S. Frota, A. Letanneux, M. Vigarario, M. Cruz, I. P. Martins, F. Viallet, and J. J. Ferreira, "Dysarthria in individuals with Parkinson's disease: a protocol for a binational, cross-sectional, case-controlled study in French and European Portuguese (FraLusoPark)," *BMJ Open*, vol. 6, no. 11, 2016.
- [17] A. Pompili, A. Abad, P. Romano, I. P. Martins, R. Cardoso, H. Santos, J. Carvalho, I. Guimaraes, and J. J. Ferreira, "Automatic detection of Parkinson's disease: an experimental analysis of common speech production tasks used for diagnosis," in *International Conference on Text, Speech, and Dialogue*. Springer, 2017, pp. 411–419.
- [18] T. Bocklet, S. Steidl, E. Nöth, and S. Skodda, "Automatic evaluation of Parkinson's speech-acoustic, prosodic and voice related cues," in *Interspeech*, 2013, pp. 1149–1153.
- [19] A. Goberman, C. Coelho, and M. Robb, "Phonatory characteristics of parkinsonian speech before and after morning medication: the ON and OFF states," *Journal of Communication Disorders*, vol. 35, no. 3, pp. 217–239, 2002.
- [20] J. Jacobi, T. Rebernik, R. Jonkers, M. Proctor, B. Maassen, and M. Wieling, "The Effect of Levodopa on Vowel Articulation in Parkinson's Disease: A Cross-Linguistic Study," in *19th International Congress of Phonetic Sciences*. Australasian Speech Science and Technology Association (ASSTA), 2019, pp. 1069–1073.
- [21] C. Moreau and S. Pinto, "Misconceptions about speech impairment in Parkinson's disease," *Movement Disorders*, vol. 34, no. 10, pp. 1471–1475, 2019.
- [22] D. A. Reynolds, T. F. Quatieri, and R. B. Dunn, "Speaker verification using adapted gaussian mixture models," *Digital Signal Processing*, vol. 10, no. 1–3, pp. 19–41, 2000.
- [23] S. Young, G. Evermann, M. Gales, T. Hain, D. Kershaw, X. Liu, G. Moore, J. Odell, D. Ollason, D. Povey, A. Ragni, V. Valtchev, P. Woodland, and C. Zhang, *The HTK Book (version 3.5a)*, Dec. 2015.
- [24] F. Eyben, K. R. Scherer, B. W. Schuller, J. Sundberg, E. André, C. Busso, L. Y. Devillers, J. Epps, P. Laukka, S. S. Narayanan, and K. P. Truong, "The Geneva Minimalistic Acoustic Parameter Set (GeMAPS) for Voice Research and Affective Computing," *IEEE Transactions on Affective Computing*, vol. 7, no. 2, pp. 190–202, 2016.
- [25] F. Eyben, F. Weninger, F. Gross, and B. Schuller, "Recent Developments in openSMILE, the Munich Open-Source Multimedia Feature Extractor," in *Proc. of the 21st ACM International Conference on Multimedia*, Barcelona, Spain, Oct. 2013, pp. 835–838.
- [26] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, software available from tensorflow.org. [Online]. Available: <http://tensorflow.org/>