



Surfboard: Audio Feature Extraction for Modern Machine Learning

Raphael Lenain, Jack Weston, Abhishek Shivkumar, Emil Fristed

Novoic Ltd

{raphael, jack, abhishek, emil}@novoic.com

Abstract

We introduce Surfboard, an open-source Python library for extracting audio features with application to the medical domain. Surfboard is written with the aim of addressing pain points of existing libraries and facilitating joint use with modern machine learning frameworks. The package can be accessed both programmatically in Python and via its command line interface, allowing it to be easily integrated within machine learning workflows. It builds on state-of-the-art audio analysis packages and offers multiprocessing support for processing large workloads. We review similar frameworks and describe Surfboard’s architecture, including the clinical motivation for its features. Using the mPower dataset, we illustrate Surfboard’s application to a Parkinson’s disease classification task, highlighting common pitfalls in existing research. The source code is opened up to the research community to facilitate future audio research in the clinical domain.

Index Terms: Audio processing, healthcare, machine learning (ML), mPower, Novoic, Parkinson’s disease, signal processing, speech and language disorders, speech representations, Surfboard.

1. Introduction

The diversity of applications of acoustic analysis is best demonstrated by the last 10 years of the INTERSPEECH Computational Paralinguistics Challenges¹, encompassing emotion detection [1], gender prediction [2], speaker state trait prediction [3, 4] and detection of medical conditions [5, 6] to name a few. Acoustic speech changes have been identified in a multitude of motor disorders (e.g. Parkinson’s disease) [7, 8, 9], affective disorders (e.g. depression) [10] and respiratory diseases (e.g. pneumonia) [11, 12]. These acoustic changes can often be detected using features extracted from speech and it is common for papers with a clinical application to define their own feature sets [7, 8, 12]. However, the selection of features and how to extract them is inconsistent across the field, resulting in a need for harmonization.

Surfboard is a Python package for audio feature extraction, written with the aim of making a library better suited to fast prototyping and modern machine learning (ML) applications than what is offered today. Our work is most similar to OpenSMILE [13], an audio feature extractor implemented in C++ that was first released in 2010 and had its latest release in 2016. OpenSMILE extracts ‘low-level descriptors’ (LLDs) from audio signals and combines them with ‘functionals’, functions that operate on time series data to extract time-independent features. Examples of LLDs include the mel-frequency cepstrum coefficients (MFCCs) and the loudness; statistics include mean and standard deviation over time.

OpenSMILE is computationally efficient but custom configuration is complex. There exists no simple Python inter-

¹<http://www.compare.openaudio.eu/>

face for OpenSMILE, hindering use in conjunction with modern ML frameworks such as scikit-learn [14], TensorFlow [15] and PyTorch [16]. Praat [17] is another popular audio feature extractor, first released in 1991. Praat is desktop-based software which since 2018 has been complemented by a Python wrapper, Parselmouth [18], and suits detailed analysis of small numbers of audio files. MATLAB [19] is frequently used by members of the audio community to extract features from speech and music, for example using Voicebox [20], MIRtoolbox [21] or, more recently, Audio Toolbox. While these are reliable toolkits, working within a MATLAB environment adds an unnecessary constraint to audio feature extraction and inhibits usability. In designing Surfboard, we attempt to combine the best of all these approaches to suit multiple use cases, including large-scale deployment.

In this paper, we first describe Surfboard’s architecture, interface, audio features and the clinical rationale behind the features. We compare these with features common to both OpenSMILE and Praat. Finally, we present a ML classifier trained on Surfboard features extracted from part of the mPower dataset [22] to highlight issues with prior work in Parkinson’s, and provide a list of reference values derived from the LibriSpeech dataset [23]. We release the Surfboard codebase² to the research community under an open-source license, along with notebooks containing all the code used in this paper³.

2. Surfboard Architecture

2.1. Overview

Surfboard aims to address the flaws of comparable frameworks while retaining their qualities. Specifically, we designed Surfboard with a focus on:

- Ease of use within Python, the lingua franca of data scientists and ML engineers [41] and the primary language for ML frameworks such as PyTorch and TensorFlow.
- The ability to process large datasets, often needed for modern ML approaches to audio processing.

Like OpenSMILE, Surfboard extracts ‘components’ (analogous to LLDs) as single values (e.g. loudness) or time series (e.g. MFCCs); in the latter case, statistics (e.g. the standard deviation) can be extracted from the time series to create time-independent features. These features can then be fed into e.g. a multilayer perceptron (MLP), a common use case for audio classification tasks. One can also obtain the full time series without extracting statistics for downstream sequential processing, for example using LSTMs [42].

The audio features currently included in Surfboard are shown in Table 1. Inspired by some of the excellent work done by the audio community, Surfboard was built on top of a number packages that we found to be well-maintained, such as

²<https://github.com/novoic/surfboard>

³<https://github.com/novoic/surfboard-IS2020>

Table 1: Description of the Surfboard features, including implementation, reference values and clinical rationale. Surfboard v0.1 reference values based on a 40-hour subset of LibriSpeech [23] are provided. The right half of the table is adapted from [24, 25, 26, 10] and summarizes clinical validation of recent review papers across indications. \uparrow = feature increases compared with healthy controls; \downarrow = feature decreases compared with healthy controls; \updownarrow = feature can increase or decrease compared with healthy controls, depending on derived feature (e.g. which MFCC component). \leftrightarrow is used to indicate that features have been applied for classification, but that how they change is unknown. For WhC, the symbol \leftrightarrow has been used: none of the reviewed papers on respiratory conditions reported feature values but the inclusion of spectral features was indeed motivated by the respiratory literature. - = unknown. PD = Parkinson’s disease; MND = motor neurone disease, synonymous with amyotrophic lateral sclerosis (ALS); MS = multiple sclerosis; HD = Huntington’s disease; MDD = major depressive disorder; HpM = hypomania; Anx = anxiety; Szo = schizophrenia; PTSD = post-traumatic stress disorder; WhC = whooping cough, synonymous with pertussis. A dagger (\dagger) indicates that the feature is a time series, so the reference value shown is the mean over time. * The pitch period entropy (PPE) reference method [27] sometimes produces large negative outliers and the crest factor method [28] large positive outliers, so we show here the more meaningful median and the median absolute deviation statistics in lieu of the mean and standard deviation; we also note that PPE was developed to assess sustained phonations rather than free speech. Novel, robust implementations of these features will be added in a future version of Surfboard.

Component	Impl.	LibriSpeech	PD	MND	MS	HD	MDD	HpM	Anx	Szo	PTSD	WhC
Entire waveform representations												
MFCCs	LibROSA	-	-	$\dagger[29, 8, 30]$	-	$\dagger[31]$	$\downarrow[10]$	$\downarrow[10]$	$\downarrow[10]$	-	-	$\leftrightarrow[12]$
Log mel spectrogram	LibROSA	-	-	-	-	-	-	-	-	-	-	-
Morlet continuous wavelet transform	SciPy [32]	-	-	-	-	-	-	-	-	-	-	-
Bark spectrogram	Ours	-	-	-	-	-	-	-	-	-	-	-
Magnitude spectrum	LibROSA	-	-	-	-	-	-	-	-	-	-	-
Chromas (music motivated)												
Chromagram with STFT	LibROSA	-	-	-	-	-	-	-	-	-	-	-
Chromagram with CQT	LibROSA	-	-	-	-	-	-	-	-	-	-	-
Chroma CENS	LibROSA	-	-	-	-	-	-	-	-	-	-	-
Spectral features												
Spectral slope \dagger	Ours	$(-1.10 \pm 0.412) \times 10^{-3}$	-	-	-	-	-	-	-	-	-	$\leftrightarrow[12]$
Spectral flux \dagger	Ours	$(15.2 \pm 5.64) \times 10^{-3}$	-	-	-	-	-	-	-	-	-	-
Spectral entropy \dagger	Ours	4.46 \pm 0.352	-	-	-	-	-	-	-	-	-	-
Spectral centroid \dagger	Ours	$(1.70 \pm 0.401) \times 10^3$ Hz	-	-	-	-	-	-	-	-	-	-
Spectral spread \dagger	Ours	$(1.50 \pm 0.178) \times 10^3$ Hz	-	-	-	-	-	-	-	-	-	$\leftrightarrow[12]$
Spectral skewness \dagger	Ours	$(1.74 \pm 0.621) \times 10^{-3}$	-	-	-	-	-	-	-	-	-	$\leftrightarrow[12]$
Spectral kurtosis \dagger	Ours	-2.99 \pm 0.00443	-	-	-	-	-	-	-	-	-	$\leftrightarrow[12]$
Spectral flatness \dagger	LibROSA	$(1.86 \pm 15.4) \times 10^{-3}$	-	-	-	-	-	-	-	-	-	$\leftrightarrow[12]$
Spectral rolloff \dagger	LibROSA	$(3.13 \pm 0.677) \times 10^3$ Hz	-	-	-	-	-	-	-	-	-	$\leftrightarrow[12]$
Classical speech features												
F0 contour \dagger	pysptk	149 \pm 35.6 Hz	$\dagger[30]$	$\dagger[33]$	$\dagger[26]$	$\downarrow[31, 25]$	$\downarrow[10]$	$\uparrow[10]$	$\uparrow[10]$	$\uparrow[10]$	-	-
F0 SD	pysptk	26.5 \pm 10.7 Hz	$\downarrow[24, 34, 30]$	$\downarrow[33]$	$\downarrow[26]$	$\uparrow[31, 25]$	$\uparrow[10]$	-	$\uparrow[10]$	$\downarrow[10]$	$\downarrow[10]$	-
Intensity \dagger	Ours	$(4.16 \pm 5.63) \times 10^{-3}$	-	-	-	-	$\downarrow[10]$	-	$\downarrow[10]$	$\downarrow[10]$	-	-
Intensity SD	Ours	$(6.33 \pm 5.61) \times 10^{-3}$	$\downarrow[24]$	-	-	-	$\downarrow[10]$	-	-	$\uparrow[10]$	-	-
Sliding-window root mean square (energy) \dagger	LibROSA	0.0444 \pm 0.0201	$\dagger[24]$	-	$\downarrow[26]$	-	-	-	-	-	-	-
Log energy	Ours	-25.0 \pm 3.22	$\dagger[29]$	-	-	-	-	-	-	-	-	-
Sliding-window log energy \dagger	Ours	-34.7 \pm 4.81	-	-	-	-	-	-	-	-	-	-
Zero-crossing rate	LibROSA	0.0528 \pm 0.0183	-	-	-	-	-	-	-	-	-	$\leftrightarrow[12]$
Sliding-window zero-crossing rate \dagger	LibROSA	0.0527 \pm 0.0182	-	-	-	-	-	-	-	-	-	-
Number of zero-crossings	LibROSA	$(2.92 \pm 1.34) \times 10^4$	-	-	-	-	-	-	-	-	-	-
Loudness	pyloudnorm [35]	-24.5 \pm 2.89 dB	$\downarrow[36]$	$\downarrow[33]$	$\downarrow[26]$	-	-	-	-	-	-	-
Loudness variation (sliding-window SD)	pyloudnorm	5.80 \pm 2.72 dB	$\downarrow[36]$	$\downarrow[33]$	$\uparrow[26]$	-	-	-	-	-	-	-
Crest factor* \dagger	Ours	4.35 \pm 1.15 (median \pm MAD)	-	-	-	-	-	-	-	-	-	-
Motivated by the clinical literature												
Pitch period entropy*	Ours	3.96 \pm 3.37 (median \pm MAD)	$\uparrow[7]$	-	-	$\uparrow[31]$	-	-	-	-	-	-
Jitter variants	-	-	$\dagger[30]$	$\dagger[33, 37]$	$\dagger[26, 38, 39]$	$\dagger[31, 25]$	$\uparrow[10]$	-	$\uparrow[10]$	-	-	-
Jitter (local)	Ours	0.0128 \pm 0.00374	-	-	-	-	-	-	-	-	-	-
Jitter (local, absolute)	Ours	$(9.31 \pm 2.97) \times 10^{-5}$ s	-	-	-	-	-	-	-	-	-	-
Jitter (RAP)	Ours	$(3.14 \pm 0.928) \times 10^{-3}$	-	-	-	-	-	-	-	-	-	-
Jitter (PPQ5)	Ours	$(5.53 \pm 1.62) \times 10^{-3}$	-	-	-	-	-	-	-	-	-	-
Jitter (DDP)	Ours	$(9.43 \pm 2.78) \times 10^{-3}$	-	-	-	-	-	-	-	-	-	-
Shimmer variants	-	-	$\dagger[8, 30]$	-	$\dagger[26, 38, 39]$	$\dagger[31, 25]$	$\uparrow[10]$	-	$\uparrow[10]$	-	-	-
Shimmer (local)	Ours	0.0966 \pm 0.0231	-	-	-	-	-	-	-	-	-	-
Shimmer (local, db)	Ours	0.737 \pm 0.113 dB	-	-	-	-	-	-	-	-	-	-
Shimmer (APQ3)	Ours	0.0363 \pm 0.00906	-	-	-	-	-	-	-	-	-	-
Shimmer (APQ5)	Ours	0.0615 \pm 0.0161	-	-	-	-	-	-	-	-	-	-
Shimmer (APQ11)	Ours	0.135 \pm 0.0497	-	-	-	-	-	-	-	-	-	-
Detrended fluctuation analysis	Ours	0.940 \pm 0.152	$\dagger[7, 8, 29]$	-	-	$\dagger[31, 31]$	-	-	-	-	-	-
Linear spectral coefficients	LibROSA	-	-	-	-	-	-	$\uparrow[10]$	-	$\uparrow[10]$	-	-
Linear spectral frequencies	Ours	-	-	-	-	-	-	-	-	-	-	-
Formant F1	Ours	$(1.16 \pm 0.455) \times 10^3$ Hz	$\dagger[40]$	-	-	-	$\downarrow[10]$	$\uparrow[10]$	$\dagger[10]$	$\dagger[10]$	-	-
Formant F2	Ours	$(1.93 \pm 0.468) \times 10^3$ Hz	-	-	-	-	-	-	-	-	-	-
Formant F3	Ours	$(2.73 \pm 0.452) \times 10^3$ Hz	-	-	-	-	-	-	-	-	-	-
Formant F4	Ours	$(3.52 \pm 0.453) \times 10^3$ Hz	-	-	-	-	-	-	-	-	-	-
Formant $\Delta F1$ \dagger	Ours	0.0417 \pm 1.17	-	-	-	-	-	-	-	-	-	-
Formant $\Delta F2$ \dagger	Ours	0.0564 \pm 1.12	-	$\downarrow[33]$	$\downarrow[26]$	-	-	-	-	-	-	-
Formant $\Delta F3$ \dagger	Ours	0.0883 \pm 1.15	-	-	-	-	-	-	-	-	-	-
Formant $\Delta F4$ \dagger	Ours	0.123 \pm 1.21	-	-	-	-	-	-	-	-	-	-
Sliding-window formant F1 \dagger	Ours	$(1.27 \pm 0.150) \times 10^3$ Hz	-	-	-	-	-	-	-	-	-	-
Sliding-window formant F2 \dagger	Ours	$(2.20 \pm 0.140) \times 10^3$ Hz	-	-	-	-	-	-	-	-	-	-
Sliding-window formant F3 \dagger	Ours	$(3.09 \pm 0.131) \times 10^3$ Hz	-	-	-	-	-	-	-	-	-	-
Sliding-window formant F4 \dagger	Ours	$(3.96 \pm 0.137) \times 10^3$ Hz	-	-	-	-	-	-	-	-	-	-
Sliding-window amplitude kurtosis \dagger	Ours	1.74 \pm 1.96	-	-	-	-	-	-	-	-	-	-
Amplitude shannon entropy	Ours	$(8.28 \pm 6.93) \times 10^3$	-	-	-	-	-	-	-	-	-	-
HNR	Ours	9.11 \pm 2.29 dB	$\downarrow[7, 8]$	-	-	$\downarrow[31, 25]$	$\uparrow[10]$	-	-	-	-	-

LibROSA [43] and pysptk [44]. Further external implementations are referenced in Table 1. We first picked components and statistics which we deemed the most prominent in the ML audio/speech analysis literature, for example those in [45, 46, 47]. We then reviewed similar frameworks, including OpenSMILE

and Praat, and conducted a clinical literature review of the application of speech to medical diagnosis to identify further features for inclusion. In [29, 8, 30, 48] for instance, the authors make use of features including the jitter variants, the shimmer variants and the harmonics-to-noise ratio to detect Parkinson’s

disease (PD) from speech. Another example from [12] is the crest factor, which can be used for whooping cough detection. More details of motivation can be found in Table 1.

Unlike OpenSMILE, Surfboard is released under an open-source license (the GNU GPL v3). We do this to ensure that the research community has the freedom to use and modify this software as they please, to empower new open source libraries built using Surfboard, and in the hope of fostering an active community of contributors.

2.2. Using Surfboard

There are two ways to use the package:

- **Native Python:** The user imports the Surfboard module and instantiates the `Waveform` class to load an audio signal from an array or a file. Features are then extracted by calling `Waveform`'s methods, or in batches using the `extract_features` helper function. This mode of use was designed for data exploration, medium-scale and on-the-fly feature extraction, for example when training a moderately sized ML model. The output is a `Pandas DataFrame` with each row representing the features extracted from a single waveform.
- **Command line interface (CLI):** The CLI is designed to extract features from a folder of audio files, given a configuration YAML file describing the desired features (i.e. the combination of components and statistics). The output is a CSV file corresponding to the `DataFrame` described above. See the documentation in our codebase for more details. We designed the CLI envisioning use for large-scale feature extraction requiring multiprocessing and/or submission to cloud virtual machine instances or local clusters.

2.3. Feature Comparison with OpenSMILE and Praat

Surfboard, Praat and OpenSMILE extract different features, with a significant overlap. We consider a subset of the features (local shimmer, local jitter, DDP jitter) offered by all three frameworks. None of the three produces directly comparable values using default parameters due to substantial implementation differences. We therefore compare them using Spearman's rank correlation coefficient. The values obtained are given in Table 2.

Table 2: Comparison of the three frameworks using Spearman's rank correlation coefficient. *LJ* = local jitter; *DDPJ* = DDP jitter; *LS* = local shimmer.

Comparison	ρ_{LJ}	ρ_{DDPJ}	ρ_{LS}
Praat vs OpenSMILE	0.30	0.31	0.43
Surfboard vs OpenSMILE	0.21	0.18	0.40
Surfboard vs Praat	0.32	0.13	0.33

We note significant rank correlation between the chosen features across all three pairs of frameworks. However, none of the three pairs agree perfectly. This can be attributed to differing implementations; for instance, Surfboard makes use of the RAPT pitch-tracking algorithm [49] inspired by [29], whereas OpenSMILE and Praat each employ custom peak-picking algorithms (see [17] and the OpenSMILE codebase⁴) for jitter/shim-

⁴<https://www.audeering.com/opensmile/>

mer calculations. Furthermore, feature extraction functions typically require parameters (e.g. sampling rate of the waveform and hop length); here we choose framework-dependent default parameters, which could impact the rank correlations.

3. Application: Classifying Parkinson's Disease Using the mPower Dataset

3.1. Experimental Design

A substantial number of the components developed in the Surfboard package were motivated by the clinical literature. Speech changes have been reported in a multitude of diseases (see Table 1); in this section we take the example of Parkinson's disease (PD). The main symptoms of PD include tremor and rigidity but effects on the motor system extend to the vocal cord, where vocal impairments are common [50, 51], with up to 70-90% prevalence after the onset of the disease [52, 53]. These vocal impairments may be one of the earliest indicators of disease [54] and deterioration of speech accompanies PD progression [54]. These early signs include reduced voice volume (hypophonia) and breathiness, hoarseness or creakiness in the voice (dysphonia), preceding more generalized speech disorder [52, 53]. Such impairments can be detected using audio analysis; in [7], the authors collected sustained phonations from 31 participants and selected 10 measures to use as input to a classifier, achieving 91.4% accuracy. A follow-up study extended this set of features and achieved 99% accuracy [8]. However, [9] criticized their work by arguing that the dataset was small (263 phonations) and that their training and test sets featured the same participants, an antipattern in ML that can lead to scientifically invalid conclusions.

We illustrate the application of Surfboard to the task of classifying PD sufferers versus healthy controls (HC). We work with the mPower dataset, a large, real-world dataset of sustained phonations from PD sufferers and HCs. Illustrating how test set design can be leveraged to achieve more rigorous evaluation metrics, we first create a subset S of 12,094 of the 62,609 phonations, leaving 50,515 for training. We ensure that S is balanced in terms of diagnosis (6,157 labelled PD and 6,157 labelled HC) and further split S into three test sets, S_1 , S_2 and S_3 , such that:

- S_1 comprises 2,500 phonations labelled HC and 2,500 phonations labelled PD, randomly sampled from S . We do not restrict the inclusion of phonations from participants already present in the training set.
- S_2 comprises 2,547 phonations labelled HC and 2,547 phonations labelled PD, randomly sampled from $S \setminus S_1$ with the additional constraint that the phonations contained in S_2 cannot be produced by participants already in the training set (i.e. S_2 and the training set are disjoint by participant).
- S_3 comprises the remaining 1,000 HC participants and 1,000 PD participants. We ensure that S_3 is age- and gender-matched, and that S_3 and the training set are disjoint by participant.

We extract features using Surfboard v0.1, choosing the subset of Surfboard components with demonstrated clinical relevance to PD (see Table 1): MFCCs, jitter variants, shimmer variants, PPE, HNR, loudness, formants, log energy, RMS. We use the entire statistics set offered by Surfboard to generate feature vectors from the time series components, resulting in one 377-dimensional vector per phonation. If feature extraction

fails, for example if the F0 extraction fails to recognize voiced segments for a given phonation, we replace missing fields with column averages, as is common with tabular data (albeit flawed [55]).

We use the remaining 50,515 phonations to train a single gradient boosting classifier using the scikit-learn library and evaluate separately on S_1 , S_2 and S_3 . We hypothesize that a decrease would be seen in our trained classifier’s performance between S_1 and S_2 , since the classifier can no longer benefit from merely identifying participants. We likewise expect a deterioration in performance between S_2 and S_3 , since the classifier can no longer learn to leverage the difference in age and gender distributions between the PD and HC subsets of the data.

3.2. Results

The results are shown in Table 3, illustrating a progressive decrease in performance. This supports our hypothesis and the criticism raised in [9]. In this work, we do not carry out a thorough feature selection process to optimize classification metrics. Instead, we illustrate how Surfboard can be used for audio analysis and exhibit some of the prior flaws in the existing literature, hopefully stimulating additional statistical rigor in future PD research.

Table 3: PD/HC classification results on the three test sets. See main text for a description of S_1 , S_2 and S_3 .

Test set	Precision	Recall	Accuracy	AUC
S_1	0.82	0.80	0.80	0.90
S_2	0.74	0.73	0.73	0.79
S_3	0.72	0.70	0.70	0.78

4. Conclusion

This paper presented Surfboard, a Python package for clinical audio analysis using modern ML workflows. We described Surfboard’s high-level architecture and the rationale behind our design choices. A comparison of features implemented in Surfboard, Praat and OpenSMILE showed that all three differ substantially, even using a measure of rank correlation in lieu of comparing absolute values. This result emphasizes the need for harmonization of speech features to facilitate more consistent research. The clinical significance of Surfboard features was demonstrated through a thorough literature review and the development of a state-of-the-art Parkinson’s disease classifier using voice data from the mPower dataset, which also addressed persistent issues of bias in existing work in this area. Our hope is that Surfboard will provide this much-needed harmonization and inspire more consistent research in the field of audio analysis, particularly at its intersection with the clinical domain.

5. Acknowledgements

We would like to acknowledge and thank the participants of the mPower study who made our use case of classifying PD from sustained phonations possible. This data was contributed by users of the Parkinson’s mPower mobile application as part of the mPower study. The study was run by Sage Bionetworks and data access is managed via the Synapse platform⁵.

⁵<https://www.synapse.org/>

6. References

- [1] B. Schuller, S. Steidl, and A. Batliner, “The interspeech 2009 emotion challenge,” in *Tenth Annual Conference of the International Speech Communication Association*, 2009.
- [2] B. Schuller, S. Steidl, A. Batliner, F. Burkhardt, L. Devillers, C. Muller, and S. S. Narayanan, “The interspeech 2010 paralinguistic challenge,” in *Eleventh Annual Conference of the International Speech Communication Association*, 2010.
- [3] B. Schuller, S. Steidl, A. Batliner, F. Schiel, and J. Krajewski, “The interspeech 2011 speaker state challenge,” in *Twelfth Annual Conference of the International Speech Communication Association*, 2011.
- [4] B. Schuller, S. Steidl, A. Batliner, E. Noth, A. Vinciarelli, F. Burkhardt, R. v. Son, F. Weninger, F. Eyben, T. Bocklet *et al.*, “The interspeech 2012 speaker trait challenge,” in *Thirteenth Annual Conference of the International Speech Communication Association*, 2012.
- [5] B. Schuller, S. Steidl, A. Batliner, A. Vinciarelli, K. Scherer, F. Ringeval, M. Chetouani, F. Weninger, F. Eyben, E. Marchi *et al.*, “The interspeech 2013 computational paralinguistics challenge: Social signals, conflict, emotion, autism,” in *Proceedings INTERSPEECH 2013, 14th Annual Conference of the International Speech Communication Association, Lyon, France*, 2013.
- [6] B. Schuller, S. Steidl, A. Batliner, S. Hantke, F. Honig, J. R. Orozco-Arroyave, E. Noth, Y. Zhang, and F. Weninger, “The interspeech 2015 computational paralinguistics challenge: nativeness, parkinson’s & eating condition,” in *Sixteenth annual conference of the international speech communication association*, 2015.
- [7] M. Little, P. McSharry, E. Hunter, J. Spielman, and L. Ramig, “Suitability of dysphonia measurements for telemonitoring of parkinsons disease,” *Nature Precedings*, pp. 1–1, 2008.
- [8] A. Tsanas, M. A. Little, P. E. McSharry, J. Spielman, and L. O. Ramig, “Novel speech signal processing algorithms for high-accuracy classification of parkinson’s disease,” *IEEE transactions on biomedical engineering*, vol. 59, no. 5, pp. 1264–1271, 2012.
- [9] J. Orozco-Arroyave, F. Honig, J. Arias-Londoño, J. Vargas-Bonilla, K. Daqrouq, S. Skodda, J. Ruzs, and E. Noth, “Automatic detection of parkinson’s disease in running speech spoken in three different languages,” *The Journal of the Acoustical Society of America*, vol. 139, no. 1, pp. 481–500, 2016.
- [10] D. M. Low, K. H. Bentley, and S. S. Ghosh, “Automated assessment of psychiatric disorders using speech: A systematic review,” *Laryngoscope Investigative Otolaryngology*, vol. 5, no. 1, pp. 96–116, 2020.
- [11] U. R. Abeyratne, V. Swarnkar, A. Setyati, and R. Triasih, “Cough sound analysis can rapidly diagnose childhood pneumonia,” *Annals of biomedical engineering*, vol. 41, no. 11, pp. 2448–2462, 2013.
- [12] R. X. A. Pramono, S. A. Imtiaz, and E. Rodriguez-Villegas, “A cough-based algorithm for automatic diagnosis of pertussis,” *PloS one*, vol. 11, no. 9, 2016.
- [13] F. Eyben, M. Wollmer, and B. Schuller, “Opensmile: the munich versatile and fast open-source audio feature extractor,” in *Proceedings of the 18th ACM international conference on Multimedia*, 2010, pp. 1459–1462.
- [14] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg *et al.*, “Scikit-learn: Machine learning in python,” *Journal of machine learning research*, vol. 12, no. Oct, pp. 2825–2830, 2011.
- [15] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard *et al.*, “Tensorflow: A system for large-scale machine learning,” in *12th Symposium on Operating Systems Design and Implementation*, 2016, pp. 265–283.

- [16] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems*, 2019, pp. 8024–8035.
- [17] P. Boersma and D. Weenink, "Praat: Doing phonetics by computer [computer program]. version 6.0.37," *Retrieved February*, vol. 3, p. 2018, 2018.
- [18] Y. Jadoul, B. Thompson, and B. De Boer, "Introducing parselmouth: A python interface to praat," *Journal of Phonetics*, vol. 71, pp. 1–15, 2018.
- [19] MATLAB, version 7.10.0 (R2010a). Natick, Massachusetts: The MathWorks Inc., 2010.
- [20] M. Brookes *et al.*, "Voicebox: Speech processing toolbox for matlab," *Software, available [Mar. 2011] from www.ee.ic.ac.uk/hp/staff/dmb/voicebox/voicebox.html*, vol. 47, 1997.
- [21] O. Lartillot and P. Toivianen, "A matlab toolbox for musical feature extraction from audio," in *International conference on digital audio effects*. Bordeaux, 2007, pp. 237–244.
- [22] B. M. Bot, C. Suver, E. C. Neto, M. Kellen, A. Klein, C. Bare, M. Doerr, A. Pratap, J. Wilbanks, E. R. Dorsey *et al.*, "The mpower study, parkinson disease mobile data collected using researchkit," *Scientific data*, vol. 3, no. 1, pp. 1–9, 2016.
- [23] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [24] V. Boschi, E. Catricala, M. Consonni, C. Chesi, A. Moro, and S. F. Cappa, "Connected speech in neurodegenerative language disorders: a review," *Frontiers in psychology*, vol. 8, p. 269, 2017.
- [25] J. C. Chan, J. C. Stout, and A. P. Vogel, "Speech in prodromal and symptomatic huntingtons disease as a model of measuring onset and progression in dominantly inherited neurodegenerative diseases," *Neuroscience & Biobehavioral Reviews*, 2019.
- [26] G. Noffs, T. Perera, S. C. Kolbe, C. J. Shanahan, F. M. Boonstra, A. Evans, H. Butzkueven, A. van der Walt, and A. P. Vogel, "What speech can tell us: A systematic review of dysarthria characteristics in multiple sclerosis," *Autoimmunity reviews*, vol. 17, no. 12, pp. 1202–1209, 2018.
- [27] M. Vashkevich. Troparion github repository. [Online]. Available: <https://github.com/Mak-Sim/Troparion/>
- [28] S. Boyd, "Multitone signals with low crest factor," *IEEE transactions on circuits and systems*, vol. 33, no. 10, pp. 1018–1022, 1986.
- [29] A. Tsanas, M. A. Little, P. E. McSharry, and L. O. Ramig, "Non-linear speech analysis algorithms mapped to a standard metric achieve clinically useful quantification of average parkinson's disease symptom severity," *Journal of the royal society interface*, vol. 8, no. 59, pp. 842–855, 2011.
- [30] T. Bocklet, S. Steidl, E. Noth, and S. Skodda, "Automatic evaluation of parkinson's speech-acoustic, prosodic and voice related cues," in *INTERSPEECH*, 2013, pp. 1149–1153.
- [31] J. Rusz, C. Saft, U. Schlegel, R. Hoffman, and S. Skodda, "Phonatory dysfunction as a preclinical symptom of huntington disease," *PLoS one*, vol. 9, no. 11, 2014.
- [32] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright *et al.*, "Scipy 1.0: fundamental algorithms for scientific computing in python," *Nature methods*, vol. 17, no. 3, pp. 261–272, 2020.
- [33] B. Tomik and R. J. Guiloff, "Dysarthria in amyotrophic lateral sclerosis: A review," *Amyotrophic Lateral Sclerosis*, vol. 11, no. 1-2, pp. 4–15, 2010.
- [34] S. Skodda, W. Gronheit, 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.
- [35] C. Steinmetz. pypstkt github repository. [Online]. Available: <https://github.com/csteinmetz1/pyloudnorm>
- [36] A. Bayestehtashk, M. Asgari, I. Shafran, and J. McNames, "Fully automated assessment of the severity of parkinson's disease from speech," *Computer speech & language*, vol. 29, no. 1, pp. 172–185, 2015.
- [37] A. K. Silbergleit, A. F. Johnson, and B. H. Jacobson, "Acoustic analysis of voice in individuals with amyotrophic lateral sclerosis and perceptually normal vocal quality," *Journal of Voice*, vol. 11, no. 2, pp. 222–231, 1997.
- [38] M. Dogan, E. Eryuksel, I. Kocak, T. Celikel, and M. A. Sehitoglu, "Subjective and objective evaluation of voice quality in patients with asthma," *Journal of Voice*, vol. 21, no. 2, pp. 224–230, 2007.
- [39] A. V. Feijó, M. A. Parente, M. Behlau, S. Haussen, M. C. De Vecchio, and B. C. de Faria Martignago, "Acoustic analysis of voice in multiple sclerosis patients," *Journal of Voice*, vol. 18, no. 3, pp. 341–347, 2004.
- [40] J. Rusz, R. Cmejla, T. Tykalova, H. Ruzickova, J. Klempir, V. Majerova, J. Picmausova, J. Roth, and E. Ruzicka, "Imprecise vowel articulation as a potential early marker of parkinson's disease: Effect of speaking task," *The Journal of the Acoustical Society of America*, vol. 134, no. 3, pp. 2171–2181, 2013.
- [41] Pypl popularity of programming language. [Online]. Available: <http://pypl.github.io/PYPL.html>
- [42] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [43] B. McFee, C. Raffel, D. Liang, D. P. Ellis, M. McVicar, E. Battenberg, and O. Nieto, "librosa: Audio and music signal analysis in python," in *Proceedings of the 14th python in science conference*, vol. 8, 2015.
- [44] R. Yamamoto, J. Felipe, and M. Blaauw. pypstkt github repository. [Online]. Available: <https://github.com/r9y9/pypstkt>
- [45] J. F. Gemmeke, L. Vuegen, P. Karsmakers, B. Vanrumste *et al.*, "An exemplar-based nmf approach to audio event detection," in *2013 IEEE workshop on applications of signal processing to audio and acoustics*. IEEE, 2013, pp. 1–4.
- [46] X. Zhuang, X. Zhou, M. A. Hasegawa-Johnson, and T. S. Huang, "Real-world acoustic event detection," *Pattern Recognition Letters*, vol. 31, no. 12, pp. 1543–1551, 2010.
- [47] A. Mesaros, T. Heittola, A. Eronen, and T. Virtanen, "Acoustic event detection in real life recordings," in *2010 18th European Signal Processing Conference*. IEEE, 2010, pp. 1267–1271.
- [48] J. P. Teixeira, C. Oliveira, and C. Lopes, "Vocal acoustic analysis-jitter, shimmer and hnr parameters," 2013.
- [49] D. Talkin and W. B. Kleijn, "A robust algorithm for pitch tracking (rapt)," *Speech coding and synthesis*, vol. 495, p. 518, 1995.
- [50] A. J. Hughes, S. E. Daniel, S. Blankson, and A. J. Lees, "A clinicopathologic study of 100 cases of parkinson's disease," *Archives of neurology*, vol. 50, no. 2, pp. 140–148, 1993.
- [51] D. G. Hanson, B. R. Gerratt, and P. H. Ward, "Cinegraphic observations of laryngeal function in parkinson's disease," *The Laryngoscope*, vol. 94, no. 3, pp. 348–353, 1984.
- [52] A. K. Ho, R. Ianseck, C. Marigliani, J. L. Bradshaw, and S. Gates, "Speech impairment in a large sample of patients with parkinsons disease," *Behavioural neurology*, vol. 11, no. 3, pp. 131–137, 1999.
- [53] 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.
- [54] B. Harel, M. Cannizzaro, and P. J. Snyder, "Variability in fundamental frequency during speech in prodromal and incipient parkinson's disease: A longitudinal case study," *Brain and cognition*, vol. 56, no. 1, pp. 24–29, 2004.
- [55] A. N. Baraldi and C. K. Enders, "An introduction to modern missing data analyses," *Journal of school psychology*, vol. 48, no. 1, pp. 5–37, 2010.