Towards robust paralinguistic assessment for real-world mobile health (mHealth) monitoring: an initial study of reverberation effects on speech

Judith Dineley[1], Ewan Carr[1], Faith Matcham[2], Johnny Downs[1], Richard Dobson[1,3], Thomas F. Quatieri[4], Nicholas Cummins[1]

[1]Institute of Psychiatry, Psychology and Neuroscience, King's College London, London, UK
[2]School of Psychology, University of Sussex, Falmer, UK
[3]Institute of Health Informatics, University College London, London, UK
[4]MIT Lincoln Laboratory, Lexington, MA, USA

judith.dineley@kcl.ac.uk, nick.cummins@kcl.ac.uk

Abstract

Speech is promising as an objective, convenient tool to monitor health remotely over time using mobile devices. Numerous paralinguistic features have been demonstrated to contain salient information related to an individual's health. However, mobile device specification and acoustic environments vary widely, risking the reliability of the extracted features. In an initial step towards quantifying these effects, we report the variability of 13 exemplar paralinguistic features commonly reported in the speech-health literature and extracted from the speech of 42 healthy volunteers recorded consecutively in rooms with low and high reverberation with one budget and two higher-end smartphones, and a condenser microphone. Our results show reverberation has a clear effect on several features, in particular voice quality markers. They point to new research directions investigating how best to record and process in-the-wild speech for reliable longitudinal health state assessment.

Index Terms: measurement error, reproducibility, acoustic environment, paralinguistics, mobile health

1. Introduction

Speech processing applied to in-the-wild smartphone recordings in health state assessment has potential as a convenient, objective monitoring tool in health care and research [1]. As a clinical application, this may break down barriers affected individuals face in accessing healthcare. Multiple studies have demonstrated the inherent value of the paralinguistic content of speech in the detection and classification of a range of health disorders e.g., [2]–[8]

However, much of this research has used one-off speech recordings made in highly controlled environments with high-specification audio equipment, with collection overseen by a trained researcher. In-the-wild longitudinal monitoring with mobile devices, in contrast, introduces multiple technical, acoustic and human factors, and associated variability in the recorded speech signal, some of which could be erroneously interpreted as related to a change in health state. These factors may include the specification of the recording hardware, ambient noise and room acoustics. Smartphone design varies greatly, resulting in varying data quality and reproducibility issues when recording on different devices [9].

Literature quantifying sources of variability in recorded speech for mobile health applications is relatively sparse. Most studies have focused on dysphonia and a small number of features, most commonly comparing the performance of mobile devices in capturing speech features to a studio microphone [11]–[14]. Studies have shown that fundamental frequency (F0) is more robust to device type and ambient noise, while voice quality features such as jitter and shimmer are heavily impacted [12], [13]. Device type has also been shown to affect the signal-to-noise ratio of recorded speech [15]. Voice quality features have also been shown, in a controlled phonetics study, to be affected by recording environment [16].

The effects of ambient noise and microphone type have been somewhat investigated, however, the effect of reverberation on speech features in-the-wild has received no specific attention in the literature, to the best of the authors' knowledge, and differences in reverberation are often audible. To begin to address this knowledge gap, our pilot measures the effect of reverberation on exemplar speech features commonly used in health-driven analyses, recording healthy volunteers in rooms with low and high reverberation, simultaneously with several devices. We chose to focus on the healthy state to create baseline data to facilitate the interpretation of reverberation-affected data from clinical cohorts in future work. We assessed features extracted with standard open-source processing tools to observe the sensitivity of the entire data pipeline to reverberation and not only recording hardware.

Our first results, that are part of a larger body of work, are reported here. We report two analyses. First, we assessed the effect of reverberation on features extracted at a suprasegmental level from recordings, a common approach for baseline machine learning experiments. Second, removing the effects of linguistic variability, we report changes in features with reverberation at the vowel level. We analyzed recordings from a benchmark condenser microphone and three smartphones to observe the effects of reverberation in each. Our aim was to investigate the implications of room reverberation for mHealth tool robustness in research and clinical practice.

2. Methodology

With no public dataset of mobile device recordings in controlled conditions with varying reverberation available, we created a new speech dataset. To facilitate rapid recruitment of a balanced cohort in the short time frame available, it was decided to not make these recordings publicly available. Speech features and basic non-identifiable, anonymized participant data will be available upon completion of the project per the conditions of the participants' consent. The recording and analysis pipeline is summarized in Figure 1.
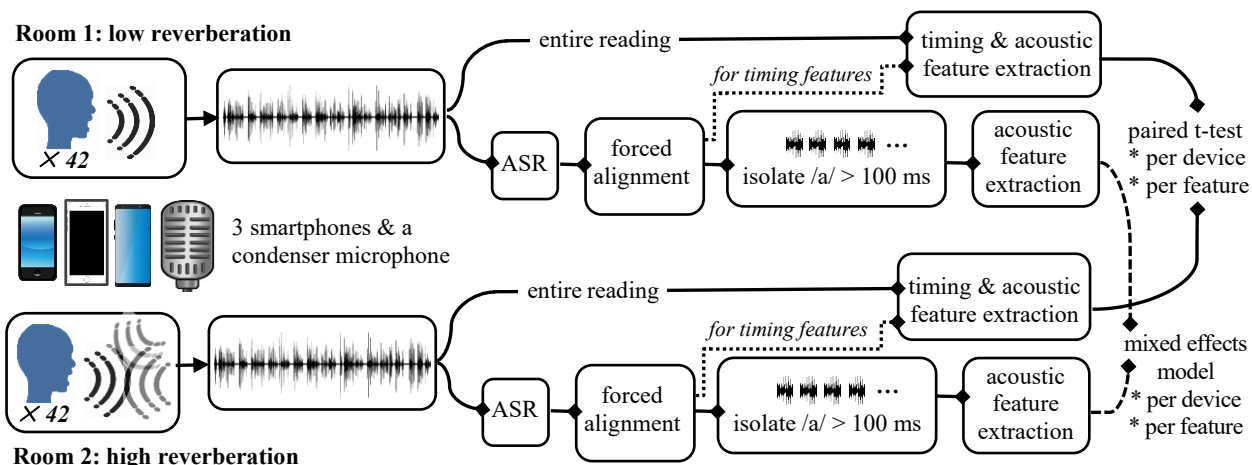
Figure 1: *Data pipeline for investigation of room reverberation effects on 13 exemplar timing and acoustic speech features extracted from smartphone recordings. Forced alignment provides information for timing feature extraction.*

2.1 Data Collection

Participant recruitment: This project received approval from the Research Ethics Committee of King's College London. Forty-two volunteers were recruited by word-of-mouth, social media and the institute's participant recruitment website and e-newsletter. Exclusion criteria included dyslexia, being a smoker, being under 16, ongoing treatment for a mental health disorder, and having any other kind of neurological, respiratory or other health issue that may affect the participant's speech. Volunteers whose first language was not English were required to have a sufficient level of reading ability and spoken English to read aloud two easy-intermediate texts and describe an everyday scene. Enrolment questionnaires recorded participants' age, sex, height, whether English was their first language and, if not, their level of English according to the Common European Framework of Reference for Languages[1]. For analysis, all participant data was saved and labelled only with a non-personally identifiable ID number.

Speech Recording: Participants were recorded in Summer 2022 in two small, neighboring, identical, windowless electroencephalogram test rooms with low ceilings in the basement of a university building where there was no ambient noise. *Room 1* was fitted with acoustic foam and soft furnishings and *Room 2* was empty, leaving only hard plaster and wooden surfaces, except for carpet tiles on the floor. Thus, we had recording spaces with low and high reverberation, respectively. In a single session, we recorded a participant consecutively in each room, alternating the room order between participants to avoid introducing bias between rooms. Participants read aloud the full version of *The Rainbow Passage* in each room [17] to elicit consistent speech across the cohort with minimal training. The reading provides variety and versatility for analysis, with high phonetic balance and full phonetic coverage, and structural and lexical complexity appropriate for our healthy, educated cohort [18], [19].

Each participant reading was recorded simultaneously with a Samsung Galaxy S20 FE 5G (released 2020), as a representative non-budget Android smartphone; a Motorola G6 Play (released 2018), as a representative budget Android smartphone; an Apple iPhone 11 (released 2019) and an Audio Technica AT2020USB+ condenser microphone with a cardioid pickup area, providing a higher quality recording from a non-mobile device within the budget constraints of our pilot study. The condenser microphone was fitted with a Rycote InVision shock mount and foam pop filter on a tabletop microphone stand and operated using Audacity open-source software (v. 3.1.3). The smartphones were positioned directly adjacent to the pop filter with their primary microphones at the height of the center of the condenser microphone diaphragm.

Participants were instructed to stand in front of and facing the pop filter. The height of the smartphones and condenser microphone was adjusted to be approximately level with the participant's mouth. Immediately prior to recording, the participant was positioned 30 cm from the condenser microphone, the distance at which its frequency response is specified. Upon completion of recording, audio files were saved and transferred to a secure cloud storage service. As with all participant data, recordings were labelled only with the participant's non-personally identifiable ID number.

2.2 Feature extraction and analysis

Feature extraction: We extracted 13 exemplar features commonly used in speech-health research that represent the basic speech production characteristics of timing, prosody, quality and articulation. These were speaking rate; articulation rate; number of pauses; mean fundamental frequency, (F0 mean); standard deviation (F0 standard deviation); intensity; first formant mean frequency, (F1 mean); second formant mean frequency, (F2 mean); spectral slope; spectral tilt; jitter; shimmer; and cepstral peak prominence, (CPP).

First, we converted all audio files to single channel 16 kHz Waveform Audio File Format (WAV) files. Audio files were then automatically transcribed, using the base Open AI Whisper model (v. 20230117) [2], choosing automated analysis over manual transcription as a more realistic solution for a real-world mHealth pipeline [20]. We aligned transcripts with the reading text with the Montreal Forced Aligner [21] and English MFA acoustic model V2.0.0a[3] to extract the three timing features for the entire reading.

Table 1: *Participant characteristics (n = 42)*

sex	female	23	English	Yes	29
	male	19	L1	No	13
age	median	28	height	median	1.70
(years)	IQR	23-32	(m)	IQR	1.63-1.80

Acoustic features were extracted at two levels: (1) suprasegmentally, at the level of the entire reading, and (2) at the vowel level, for occurrences of /a/ of at least 100 ms. We undertook a suprasegmental approach as this is a common approach in paralinguistic analyses [22] and enables the calculation of timing features. In contrast, analysis of /a/ vowel sounds removes variability due to the diverse linguistic content of the entire reading. An open vowel, /a/ has been recommended as more reliable for jitter and shimmer measurements [23], two features associated with changes in health, in particular mental state [7]. Timings provided by forced alignment enabled the extraction of feature values for each detected /a/ sound.

Acoustic features were extracted using Parselmouth (v. 0.4.1), an open-source Python library (v. 3.9.13), in combination with Praat (v. 6.3.02) [24], [25]. This extraction used default Praat settings, except for the extraction of F0, which followed the two-step approach recommended in [26]. The use of Whisper, Montreal Forced Aligner and Praat was a design choice to observe the sensitivity of feature extraction to reverberation effects using a processing pipeline of standard, open-source, well-established tools.

Statistical analysis: Tests were undertaken in R (v. 4.2.2). We analyzed the 13 features extracted at the suprasegmental level using bootstrapped paired t-tests. First, for each feature and device combination, we grouped the 84 feature values obtained for every participant in the two rooms, standardizing the values (mean = 0; standard deviation = 1). In each feature-device group, we then paired feature values per participant (Room 1, Room 2) to perform the t-tests and plotted standardized mean differences and corresponding 95% confidence intervals (CI).

We used linear mixed effects (LME) models [27] to analyze the effects of reverberation on the 10 acoustic features extracted from repeated /a/ vowel sounds identified in each of our recordings. We detected a median of 5 /a/ sounds ≥ 100 ms per recording (interquartile range: 3-10). Individual models were fitted for each feature-device combination, using groups of standardized feature values for each room-participant permutation, similar to the suprasegmental analysis. In contrast to the suprasegmental t-tests, LME models account for the clustering of repeated measures and allows us to adjust for participant gender, age, height and first language [28].

Each model considered a single speech feature as the dependent variable and included reverberation (0 = low, 1 = high), years of age, height (m), sex (0=female, male = 1), L1 English (0 = No, 1 = Yes) and room recording order (0 = Room 1 first, 1 = Room 2 first) as independent variables (Table 1). Each model included a participant random intercept to account for clustering of repeated measures. We report the β coefficients for reverberation, labelled *adjusted mean difference*, and their corresponding 95% CI. The regression coefficients represented the standard deviations difference in each speech feature between low and high reverberation rooms. Bootstrap CIs were estimated using a parametric percentile bootstrap with 500 iterations, implemented using the *confint.merMod* method from the *lme4* package for R [29].
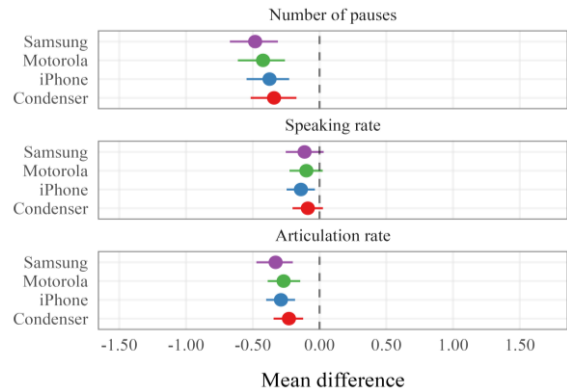

Figure 2: S*tandardized differences in timing features between low and high reverberation per participant over 42 participants. Errors bars are 95% confidence intervals. Negative differences represent lower feature values in the presence of higher reverberation.*

3. Results & Discussion

Figures 2 and 3 show the within-person standardized mean differences (and corresponding 95% confidence intervals) in each feature between the two rooms. Positive values indicate that mean feature values were higher in the presence of high reverberation (Room 2) versus low reverberation (Room 1). The three extracted timing features are all affected somewhat by room reverberation (Figure 2). The consistent negative mean difference raises robustness and validity concerns for these features, which are widely used as markers in mental health and neurological speech studies [7].

F0 mean and standard deviation are least affected by reverberation (Figure 3). Both the mean difference and room coefficient have estimates closer to zero and narrow confidence intervals. This is consistent with results reported in [12], [13] where F0 was reported as more robust to ambient noise and microphone type than other speech features. Surprisingly, we observe no clear effect of reverberation on intensity as recorded by the condenser microphone. There are, however, small offsets in intensity in all three smartphones.

Considering F1 and F2, the systematic biases observed in the suprasegmental analysis of the Motorola, iPhone and the condenser microphone are not present in the LME /a/ analysis. We speculate that this is due to the reduction in linguistic variability. Interestingly, the F1 results for Samsung differ from the other devices. We speculate that this could be due to pre-processing on the Samsung. Voice quality features appear the most affected by reverberation, with jitter, shimmer and CPP being clearly affected in all testing scenarios. These results amplify existing concerns about the robustness and validity of voice quality features for health assessment [12], [13]. We observed no clear differences between smartphones, including between the budget phone and other phones we tested. However, we plan to compare devices formally in future work.

In summary, our analysis demonstrates that room reverberation is a source of variability in multiple features. This is unsurprising; reverberation is a convolutive phenomenon and dictated by the raw speech signal and the room impulse response, and almost all extracted features are likely to be somewhat affected. A caveat is that these observations may not be due solely to reverberation; the Lombard effect may have resulted in different raw speech characteristics [30].
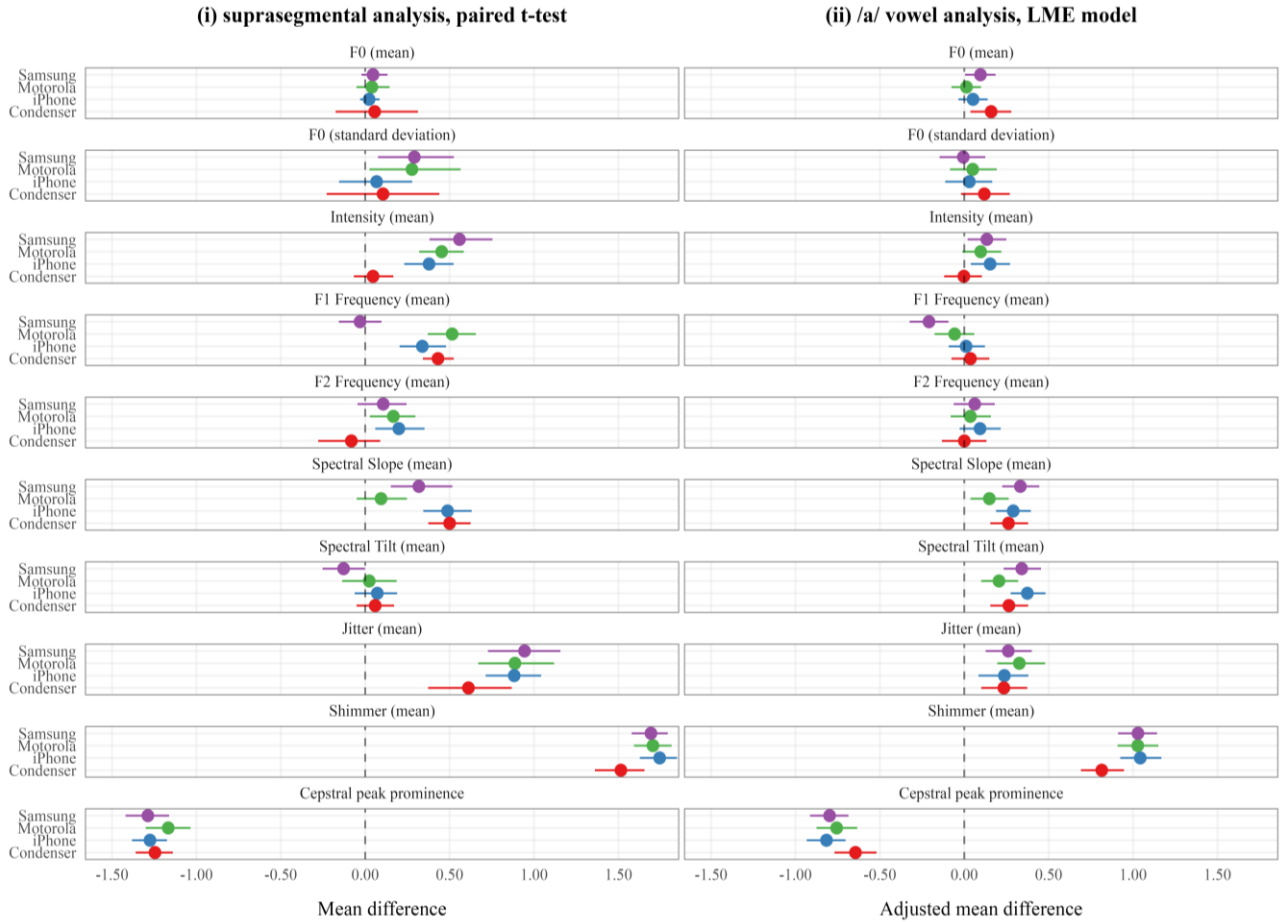
Figure 3: *Standardized differences in features values between low and high reverberation per participant over 42 participants. Mean differences in (ii) are adjusted for participant and measurement variables. Positive differences represent higher values in the presence of higher reverberation. Error bars represent 95% confidence intervals.*

Our study has several limitations. First, it was conducted in healthy volunteers. This was a design choice to begin to investigate reverberation effects while minimizing variability due to pathology and to pilot data collection procedures for future larger studies. Second, our experimental set-up likely represents the upper extreme of reverberation levels in-the-wild and consequently overestimates the effects; however, our results still highlight reliability concerns. Third, we have a limited age range. Fourth, we did not collect sustained vowels as part of our protocol; this might have allowed more reliable extraction of voice quality features. Fifth, we did not compare the effect of different processing tools and algorithm settings. Sixth, our analysis is not a formal comparison of mobile devices. We are planning a larger study to address these.

4. Conclusion

Our findings on the effects of reverberation complement observations in the literature on ambient noise and device choice [12], [13]. The speech processing community cannot ignore these sources of variability and expect to develop robust and reliable models for health assessment; otherwise, differences in recording conditions may be attributed to different health states. Instead, we must address the question of how we best record speech in-the-wild that is both robust to variations in recording parameters but still captures salient indicators of an individual's health state. Signal processing and machine learning solutions must also be explored. This could

include impulse response estimation techniques and adversarial learning approaches [31, 32]. These are vital steps in the development of transparent, reliable, valid and reproducible tools for health state assessment.

5. Acknowledgements

6. References

[1] S. Latif, J. Qadir, A. Qayyum, M. Usama, and S. Younis, "Speech Technology for Healthcare: Opportunities, Challenges, and State of the Art," *IEEE Rev. Biomed. Eng.*, vol. 14, pp. 342–356, 2021.

[2] J. R. Green et al., "Bulbar and speech motor assessment in ALS: Challenges and future directions," *Amyotroph. Lat. Scler. Frontotemp. Degener.*, vol. 14, no. 7–8, pp. 494–500, Dec. 2013.

[3] J. R. Orozco-Arroyave et al., "NeuroSpeech: An open-source software for Parkinson's speech analysis," *Digit. Sig. Proc.*, vol. 77, pp. 207–221, Jun. 2018.

[4] P. Hecker, N. Steckhan, F. Eyben, B. W. Schuller, and B. Arnrich, "Voice Analysis for Neurological Disorder Recognition–A Systematic Review and Perspective on Emerging Trends," *Front. Digit. Health*, vol. 4, p. 106, Jul. 2022.

[5] K. Vatanparvar et al., "SpeechSpiro: Lung Function Assessment from Speech Pattern as an Alternative to Spirometry for Mobile Health Tracking," in *Proc. of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*, online, Nov. 2021, pp. 7237–7243.

[6] T. Talkar, D. M. Low, A. J. Simpkin, S. Ghosh, D. T. O'Keeffe, and T. F. Quatieri, "Dissociating COVID-19 from other respiratory infections based on acoustic, motor coordination, and phonemic patterns," *Sci. Rep.* 2023 13:1, vol. 13, no. 1, pp. 1–14, Jan. 2023.

[7] D. M. Low, K. H. Bentley, and S. S. Ghosh, "Automated assessment of psychiatric disorders using speech: A systematic review," *Laryngoscope Investig. Otolaryngol.*, vol. 5, no. 1, pp. 96–116, Feb. 2020.

[8] N. Cummins, S. Scherer, J. Krajewski, S. Schnieder, J. Epps, and T. F. Quatieri, "A review of depression and suicide risk assessment using speech analysis," *Speech Commun*, vol. 71, pp. 10–49, Jul. 2015.

[9] M. Krichen, "Anomalies Detection through Smartphone Sensors: A Review," IEEE Sens. J., vol. 21, no. 6, pp. 7207–7217, Mar. 2021.

[10] P. J. Cho et al., "Demographic Imbalances Resulting From the Bring-Your-Own-Device Study Design," *JMIR Mhealth Uhealth*, vol. 10, no. 4, p. e29510, Apr. 2022.

[11] V. Uloza et al., "Exploring the feasibility of smart phone microphone for measurement of acoustic voice parameters and voice pathology screening," *Eur. Arch. Oto-Rhino-L.* 2015 272:11, vol. 272, no. 11, pp. 3391–3399, Jul. 2015.

[12] Y. Maryn, F. Ysenbaert, A. Zarowski, and R. Vanspauwen, "Mobile Communication Devices, Ambient Noise, and Acoustic Voice Measures," *J. Voice*, vol. 31, no. 2, pp. 248.e11-248.e23, Mar. 2017.

[13] S. Jannetts, F. Schaeffler, J. Beck, and S. Cowen, "Assessing voice health using smartphones: bias and random error of acoustic voice parameters captured by different smartphone types," *Int. J. Lang. Commun. Disord.*, vol. 54, no. 2, pp. 292–305, Mar. 2019.

[14] E. U. Grillo, J. N. Brosious, S. L. Sorrell, and S. Anand, "Influence of Smartphones and Software on Acoustic Voice Measures.," *Int. J. Telerehabil.*, vol. 8, no. 2, pp. 9–14, Dec. 2016.

[15] E. V. Miley, F. Schaeffler, J. Beck, M. Eichner, and S. Jannetts, "Secure account-based data capture with smartphones – preliminary results from a study of articulatory precision in clinical depression," *Linguist. Vanguard*, vol. 7, no. s1, Jan. 2021.

[16] C. Ge, Y. Xiong, and P. Mok, "How reliable are phonetic data collected remotely? Comparison of recording devices and environments on acoustic measurements," in *Proc. INTERSPEECH 2021 - 22nd Annual Conference of the International Speech Communication Association*, Brno, Czechia, Aug. 2021, pp. 3984–3988.

[17] G. Fairbanks, *Voice and Articulation Drillbook*, 2nd ed. New York: Harper, 1960.

[18] A. C. Lammert et al., "Analysis of Phonetic Balance in Standard English Passages," *J. Speech Lang. Hear. Res.*, vol. 63, no. 4, pp. 917–930, Apr. 2020.

[19] T. W. Powell, "A comparison of English reading passages for elicitation of speech samples from clinical populations," *Clin. Linguist. Phon.*, vol. 20, no. 2–3, pp. 91–97, Apr. 2009.

[20] S. Luz, F. Haider, S. De La Fuente, D. Fromm, and B. MacWhinney, "Detecting cognitive decline using speech only: The ADReSSochallenge," in *Proc. INTERSPEECH 2021 - 22nd Annual Conference of the International Speech Communication Association*, Brno, Czechia, Aug. 2021, pp. 3780–3784.

[21] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, "Montreal forced aligner: Trainable text-speech alignment using kaldi," in *Proc. INTERSPEECH 2017 - 18th Annual Conference of the International Speech Communication Association,* Stockholm, Sweden, Aug. 2017, pp. 498–502.

[22] N. Cummins, A. Baird, and B. W. Schuller, "Speech analysis for health: Current state-of-the-art and the increasing impact of deep learning," *Methods*, vol. 151, pp. 41–54, Dec. 2018.

[23] M. Brockmann, M. J. Drinnan, C. Storck, and P. N. Carding, "Reliable Jitter and Shimmer Measurements in Voice Clinics: The Relevance of Vowel, Gender, Vocal Intensity, and Fundamental Frequency Effects in a Typical Clinical Task," *J. Voice*, vol. 25, no. 1, pp. 44–53, Jan. 2011.

[24] Y. Jadoul, B. Thompson, and B. de Boer, "Introducing Parselmouth: A Python interface to Praat," *J. Phon.*, vol. 71, pp. 1–15, Nov. 2018.

[25] P. Boersma and V. Van Heuven, "PRAAT, a system for doing phonetics by computer," *Glot Int.*, vol. 5, no. 9/10, pp. 341–345, 2002.

[26] A. P. Vogel, P. Maruff, P. J. Snyder, and J. C. Mundt, "Standardization of pitch-range settings in voice acoustic analysis," *Behav. Res. Methods.*, vol. 41, no. 2, pp. 318–324, May 2009.

[27] E. Bagiella, R. P. Sloan, and D. F. Heitjan, "Mixed-effects models in psychophysiology," *Psychophysiology*, vol. 37, no. 1, pp. 13–20, Jan. 2000.

[28] B. Schuller et al., "Paralinguistics in speech and language—State-of-the-art and the challenge," *Comput. Speech Lang.*, vol. 27, no. 1, pp. 4–39, Jan. 2013.

[29] D. Bates, M. Mächler, B. M. Bolker, and S. C. Walker, "Fitting Linear Mixed-Effects Models Using lme4," *J. Stat. Softw*, vol. 67, no. 1, pp. 1–48, Oct. 2015.

[30] S. A. Zollinger and H. Brumm, "The Lombard effect," *Curr. Biol.*, vol. 21, no. 16, pp. R614–R615, Aug. 2011.

[31] G. Ciccarelli, J. Barber, A. Nair, I. Cohen, T. Zhang, "Challenges and Opportunities in Multi-device Speech Processing," in *Proc. INTERSPEECH 2022 - 23rd Annual Conference of the International Speech Communication Association,* Incheon, Korea, Sep., 2022, pp. 709–713, 2022.

[32] S. Pascual, J. Serrà, and A. Bonafonte, "Time-domain speech enhancement using generative adversarial networks," *Speech Commun.*, vol. 114, pp. 10–21, Nov. 2019.