



Visualising Model Training via Vowel Space for Text-To-Speech Systems

Binu Abeysinghe, Jesin James, Catherine I. Watson, Felix Marattukalam

Dept. of Electrical, Computer, and Software Engineering, The University of Auckland, New Zealand

babe269@aucklanduni.ac.nz, jesin.james, c.watson, felix.marattukalam@auckland.ac.nz

Abstract

With the recent developments in speech synthesis via machine learning, this study explores incorporating linguistics knowledge to visualise and evaluate synthetic speech model training. If changes to the first and second formant (in turn, the vowel space) can be seen and heard in synthetic speech, this knowledge can inform speech synthesis technology developers. A speech synthesis model trained on a large General American English database was fine-tuned into a New Zealand English voice to identify if the changes in the vowel space of synthetic speech could be seen and heard. The vowel spaces at different intervals during the fine-tuning were analysed to determine if the model learned the New Zealand English vowel space. Our findings based on vowel space analysis show that we can visualise how a speech synthesis model learns the vowel space of the database it is trained on. Perception tests confirmed that humans could perceive when a speech synthesis model has learned characteristics of the speech database it is training on. Using the vowel space as an intermediary evaluation helps understand what sounds are to be added to the training database and build speech synthesis models based on linguistics knowledge.

Index Terms: Text To Speech Synthesis, Model training, Visualisation, Vowel plots, Linguistics, Machine Learning

1. Introduction

Deep learning has risen in popularity for producing high-quality synthetic speech. Common synthetic speech evaluation include subjective tests [1, 2, 3] and objective measures [4, 3]. These evaluations are conducted *after* the speech synthesis model (hereafter called speech model) is fully trained. *During* the speech model training, we rely on the saturation of the learning curve to decide on when to stop the training. Even though the learning curve indicates how well the deep learning model learns the characteristics of the speech database it is trained on, it is insufficient to interpret the acoustic properties that the model learns and compare the properties to that of the database the model is trained on. Upon conducting subjective and objective evaluation, the approach is to re-train the model with a larger speech database if issues are found with the synthetic voice. This evaluation does not indicate which particular sounds in the language will need more representation in the database.

We propose to use the vowel space as an approach to provide an acoustic evaluation of the speech model *during* training. The vowel space is a two-dimensional area bounded by the first and second formant frequency coordinates of vowels [5]. Linguistic studies use the vowel space to study characteristics of languages [6], accents, dialects [7, 8], their changes across different regions [5], and time [9]. New language learners are trained to produce vowels as close as possible to the language's vowel space, and providing visual feedback using vowel spaces is often done [10, 11, 12, 13, 14]. This is possible because the first and second formants can be related to the jaw opening and tongue movement, respectively. Inspired by this, *we propose to use the vowel space to evaluate how a machine learning model*

learns the acoustics of the vowels of a language from a speech database. Researchers can visualise the vowel space of a given speech model and develop an understanding of how well the model learns the characteristics of the database *during* training. To develop an approach to evaluate speech model training via vowel space, we answer the following questions:

RQ1: How does the vowel space of a speech synthesis model change during the training process?

RQ2: Can humans perceive the change in vowel space of a speech synthesis model during the training process?

We train a New Zealand English (NZE) speech model to answer these questions. NZE is a distinct English accent spoken in New Zealand by 5 million people, with lexical and pronunciation differences from General American English (GAE) and British English. It is low-resourced in terms of speech technology development as there is not sufficient open-sourced data to train deep-learning-based speech models. Hence, we use a pre-trained speech model created from a large GAE database and fine-tune it using a relatively smaller NZE database. RQ1 is answered by a detailed vowel space analysis as the fine-tuning happens from GAE to NZE. RQ2 determines a link between the change in vowel space and the perception of the accent change during speech model training. A perception test and its comparison to the vowel space is conducted to answer this question.

We analyse the vowel space during speech model training, thereby providing the foundational knowledge for developing a method to visualise and evaluate speech synthesis model training. The method developed here does not replace perception tests; it provides an intermediary evaluation technique. Visualising the speech model's vowel space allows researchers to optimise against the expected vowel space of the language/speaker. This optimisation can be done via adjustments to the model parameters or using the vowel space to determine what speech sounds need to be better represented in the training database. Once the researcher is satisfied with the speech model's vowel space, they can validate the results with a perception test.

2. Background

Formants are a frequency range where there is an absolute or relative maximum in the sound spectrum. The frequency at the maximum is the formant frequency [15]. First formant, F1 corresponds to the vertical position of the tongue and is associated with frequencies between 200 and 900 Hz. Second formant, F2 corresponds to the horizontal position of the tongue and ranges from 600 - 2600 Hz [16] (See Figure 2 (b) with the NZE vowel space marked with green line). The tongue position constricts airflow in the throat and the roof of the mouth, adjusting the frequency of the air as it passes through [17]. A large part of voiced sections in speech is attributed to vowels, and vowels often exhibit a noticeable change in F1 and F2 [18]. There is strong evidence on the impact of vowel space on age [19], gender and languages [20, 9]. The vowel space is a tool to evaluate articulation problems [21, 22] and language learning [10, 14],

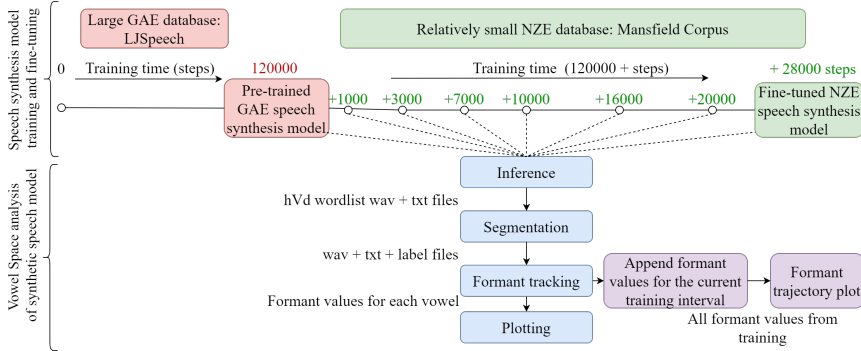


Figure 1: Flow diagram of the process of obtaining formant data when training a model

1	NZE vowels	ʌ	ɒ*	ɪ*	æ*	ʊ	e*	ɑ:*	ɜ:*	u:*	o:	i:*
2	Well’s lexical words	STRUT	LOT/CLOTH	KIT	TRAP	FOOT	DRESS	BATH/START	NURSE	GOOSE	THOUGHT	FLEECE
3	hVd word	hUd	hOd	hId	hAd	hOod	hEAd	hARd	hEARd	whO’d	hORde	hEEd
4	Word list	heard...hood...hud...heed...head...had...hard...hod...who’d...hood...heard...hid.										

Table 1: NZE vowels, usage with Well’s lexical words, hVd words and example of a word list used for perception test.

showing that it is a unique characteristic of a person’s voice for a particular language, age, gender, or speaking difficulties.

NZE, compared to GAE, is a non-rhotic English variety, sharing many features of received pronunciation. Table 1 lists the NZE vowels using International Phonetic Alphabet (Row 1) and corresponding Well’s lexical words [23] (Row 2). The NZE vowel space is marked with the green triangle in Figure 2 (b). Pertinent to this study, using Well’s lexical words, NZE has raised DRESS and TRAP vowels compared to GAE. Unlike GAE, NZE distinguishes between TRAP and BATH vowels. CLOTH vowel merges with LOT and not THOUGHT vowel in NZE. The point vowels for the NZE vowel space are FLEECE, THOUGHT, and START. The point vowels for the GAE vowel space are FLEECE, TRAP, START, and GOOSE [24].

Speech models try to learn characteristics from the database they are trained on, which in effect is learning the acoustic characteristics of the speaker(s) of the database. Although it can be argued that vowel space only captures the acoustic characteristics of the vowels, it is known that vowels are distinctive characteristics of the speaker of a language. Further, in accents of English, the main differences are due to vowel pronunciations [23, p.181].

3. Methodology

To answer RQ1, we analyse the vowel space during the training of a speech model from GAE to NZE as illustrated in Figure 1.

3.1. Speech databases

LJSpeech database [25] is the GAE database used, consisting 13,100 short audio clips of one GAE female speaker, with a duration of 24 hours. The Mansfield database [26] is the NZE database used. It consists of 1095 sentences read by a NZE female speaker with a duration of 3 hours.

3.2. Speech synthesis model training and fine-tuning

A deep learning-based speech model was trained on the LJSpeech database for 120,000 steps, as shown in Figure 1. The learning curve was observed to decide when to stop the training process. This produced the pre-trained GAE speech model. The GAE model was then fine-tuned using the Mansfield database. Again, the learning curve was observed to decide to stop the training at 28,000 additional steps. Thus, the model was trained for 148,000 steps with the last 28,000 trained exclusively on the NZE database. The vowel space analysis was conducted only on the fine-tuning to observe the changes to vowel space as the model tunes from GAE to NZE.

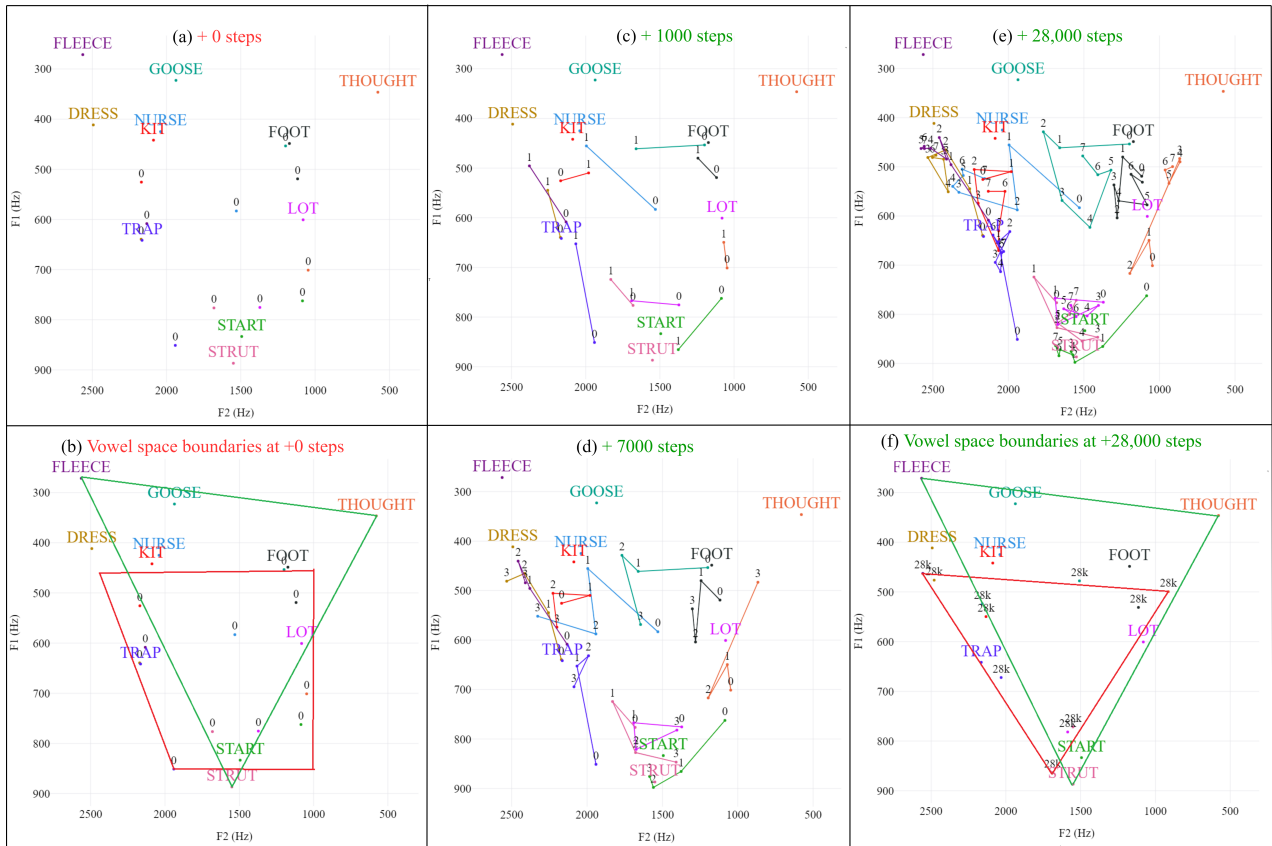
3.3. Vowel space analysis design

The model training was evaluated at multiple points (all the green steps illustrated in Figure 1) during the fine-tuning. At each step, an inference of the model trained up to that step was created. This inference was used to synthesise speech corresponding to that step. Through forced alignment, the vowel segments were identified. From these, F1, F2 were extracted and used to produce a vowel space. Each vowel point is the average of multiple renderings of the same vowel (details of sentence selection in Section 3.3.2). This vowel space represents the acoustics of vowels learned by the speech model up to that step. Figure 2(a) with points marked as 0 is the vowel space obtained at 120,000 + 0 steps - the fully GAE speech model. We also add the vowel space of a *real* NZE voice for reference [27] (produced by the same NZE speaker who produced the Mansfield database). Figure 2(c) with points marked as 1 is the vowel space at 120,000 + 1000 steps, which is the fine-tuned model after 1000 steps. The formant values obtained at the later steps are appended along with the former steps, and lines are drawn between them to indicate the vowel space progression. In Figure 2(c) the line between points marked as 0 and 1 of the same color indicate the change in vowel space as the model trained from 120,000 + 0 to 120,000 + 1000 steps. Each colour indicates each vowel in NZE, with the *real* location of the vowel represented by the Well’s lexical word of the same colour. This was then continued for the entire learning process up to 120,000 + 28000, which gives Figures 2 (d), (e).

3.3.1. Speech synthesis model and step intervals

The auto-regressive architecture Tacotron2 [28] was used to train the speech model. Tacotron2 combines two deep neural networks: Tacotron’s sequence-to-sequence spectrogram generation [29], and Google’s WaveNet [30]. Tacotron uses a Long Short Term Memory network to predict a mel spectrogram from character embeddings of the input text. The use of Tacotron in a multi-speaker scenario [31] and fine-tuning Tacotron2-based pre-trained models to that of speakers or languages with limited data [32, 33] prove that smaller datasets could benefit from transfer learning using Tacotron2. Initially, the model was trained using the LJSpeech database for 120,000 steps based on [34]. The 120,000 steps pre-trained model was then trained for a further 28,000 steps using the Mansfield database. The fine-tuning was conducted with an initial learning rate of 0.001 on Google Colab [35]. Nvidia Tesla K80 GPU was used for training, taking approximately 48 hours.

Conducting vowel space analysis at every step is time-



Vowel label : Real NZE voice vowel space Training time (steps): 0 - 0, 1 - 1000, 2 - 3000, 3 - 7000, 4 - 10000, 5 - 16000, 6 - 20000, 7 - 28000

Figure 2: The change in speech model vowel space over training steps. Each vowel point is the average of multiple vowel points.

consuming. The authors found that the most perceptual changes to the synthetic voice from GAE to NZE happened at the first 10,000 steps compared to the last 10,000 steps. Hence, more steps were chosen earlier during the fine-tuning. The steps chosen were: 0, 1000, 3000, 7000, 10,000, 16,000, 20,000, 28,000 steps; but analysis can also be conducted at all steps.

3.3.2. Sentence selection

Following linguistics studies [27], hVd words were chosen to avoid co-articulation effects (Table 1 Row 3). Five word lists were produced with the 11 hVd words in different orders to ensure that the variations in synthetic speech produced based on the neighbouring words and position of the word in a sentence were balanced out as described in [27]. Ellipses were added between each hVd word (Table 1 Row 4) to introduce a pause between them during inference. As we are fine-tuning to NZE, the analysis aims to identify how well the speech model learns NZE. Hence, the analysis considered only the NZE vowels.

3.4. Vowel space analysis method

Figure 2 shows the vowel space analysis steps. Word lists were given to the inference of fine-tuned 0 step NZE model to synthesise corresponding wav files. wav files were segmented at the phonetic level using WebMAUS [36], which has an NZE option. Text, wav file and its phonetic transcription were used to extract formants using `formant_burg()` in Praat [37] with default settings. The averages of formant values for each vowel were used to plot a vowel space using Python's `plotly` library [38]. This was repeated for selected steps to obtain Figures 2(a),(c),(d),(e), taking 40 s per step, mainly due to using external segmentation service. Based on the resources available, the analysis times can be optimised. Hand correction of segmentation and formant estimation was not done to facilitate automation.

3.5. Perception test

A perception test¹ was conducted to understand if the changes in what is learnt by the speech model *during* training are perceivable, thereby addressing RQ2. Sentences with words containing the NZE vowels along with a carrier sentence were used. The carrier sentence used was 'Say the word ... again', where a word replaced the ellipses. Based on Section 2, the vowels marked by * in Table 1 are the most different between GAE and NZE. Two words containing each vowel were chosen and placed into the carrier sentence. This resulted in 16 sentences (two each for the 8 vowels marked by *) at each of the 8 step intervals, giving a total of 128 sentences. An online perception test was conducted where the participants listened to 42 synthetic speech sentences. They rated the sentences using a Likert scale ranging from 'Not at all like this accent' to 'Completely this accent'. Along with GAE and NZE, Australian English (being similar to NZE [16]) and Canadian English (being similar to GAE [39]) were included as options. Depending on the perceived closeness to an accent, participants could move a slider on one or all the accents. 23 people participated in the survey, aged 22-57 (mean = 33, SD = 12). 13 were first language English speakers. All participants had lived in New Zealand for more than a year and considered themselves proficient in English. None of the participants reported any hearing difficulties.

4. Results and discussion

4.1. Vowel space analysis results

The vowel space analysis resulted in vowel spaces for selected steps as shown in Figure 2 (a), (c), (d), (e). It can be seen that

¹Approved by the University of Auckland Human Participants Ethics Committee (Ref.22681) on 01/10/2021 for 3 years.




Figure 3: Perception test results in a Likert scale represented using stacked bar graphs. The accent given at the top of each stacked bar plot is the accent option that was provided to participants to choose from.

there is a change in the vowel space during fine-tuning from one accent to another. The diagram illustrates that the process of vowel space transformation during speech model training is far from a linear progression. Additionally, they can be used to understand if the model can learn the vowel space accurately. We have also included the boundaries for the vowel spaces learned at 0 steps (Figure 2 (b)) and at +28,000 steps (Figure 2 (e)). Contrasting the red vowel spaces in Figure 2 (b) and (f), one can see how it changes from the GAE trapezoid shape to the NZE triangular shape, with the point vowels FLEECE, THOUGHT and START, similar to the NZE vowel space derived from the real recordings, given in green. The difference in vowel height between the two NZE vowel spaces (Figure 2 (f)) is most likely due to differences in speaking style. The real speech was spoken carefully, with hyper-articulation. The synthetic speech was produced with little pause between the words, mimicking the continuous speech it is trained on.

The figures also show how well the model can replicate the vowel space of the database it is trained on. All vowels change the most rapidly within the first 10,000 steps (Figures 2 (c), (d), (e)). Smaller variation is noticed as the model trains above 20,000 steps (Figure 2(e) compared to (d)). This is most evident in the TRAP, LOT, STRUT, FOOT vowels. These results show that the vowel space of the speech model changes towards the vowel space of the database that it is trained on, thereby answering RQ1.

4.2. Perception test results

The summary of the results for each accent option at four step intervals is shown in Figure 3². The shades of green indicate that the participants thought that the given synthetic sentence was *mainly/completely* the listed accent. At 0 steps 34%, participants thought that the speech samples were *mainly/completely* GAE. As the steps increase, the number of people who think that the speech sample given was *mainly/completely* GAE reduces (green portion decreases to 4% at 28,000 steps). Looking at the NZE result, as the number of steps increases, the number of people who perceived the speech samples as NZE also increases (20% in step 0 to 33% in step 28,000). After around 10,000 steps, the increase in identification seems to saturate. It is also important to note the large number of zero values present in each category indicates that participants generally associated the samples as one accent rather than a mix of one or more accents. Observing shades of brown at 0 steps, 45% participants thought that the GAE voice *did not sound like/slightly sounded like* GAE, and 65% of the participants thought that the NZE voice *did not sound like/slightly sounded like* NZE.

The perception test results can be linked to the vowel space

²Results of all selected steps and the other two accent options are not included in the figure due to space constraints. For the most part, Canadian English and Australian English behaved similar to GAE and NZE, respectively.

analysis. Using the vowel space analysis (Figure 2) we can see how the vowel space changes as the training progresses. From the perception test, it can be seen that the participants also perceived a change in accent from GAE to NZE as the step numbers increases. We can see that the major change in the vowel space from GAE to NZE happens in the first 10,000 steps of the training. A similar trend was observed in the perception test, where there was a sudden drop in the number of participants who perceived the synthetic speech samples as *completely/most likely* GAE at 1000 steps and an increase in participants who perceived the synthetic speech samples as *completely/most likely* NZE. Hence, RQ2 has been answered. However, the test sentence ‘Say the word ... again’ may have been insufficient to highlight the differences between vowels. More phonetically varied sentences could be chosen for detailed insights into the impact on each vowel. The process of obtaining these plots have been automated into a Python package dubbed *performant*³.

When deep learning models give inaccurate results, the solution is to train with more data or adjust training parameters. The vowel space analysis can help researchers understand which vowels need to be incorporated into the database. Also, the vowel space shape can be included as a loss function to be optimised during training. As yet, a measure for changing vowel space shape has not been automated. Also, with transformer models being used for speech synthesis, the specific vowels crucial for a language can be given more attention. Finally, the changing vowel space (Figure 2) can be used by researchers as an indicator to decide when to stop training a deep learning model, along with the learning curve which is usually used.

5. Conclusion

This study sought to explore a heavily researched aspect of linguistics - vowel spaces, and see if they were as helpful in evaluating deep learning-based synthetic speech models. This study provides researchers an easy means to visualise the closeness of a speech model to what is expected for a language, accent or speaker *during* the model training. It provides researchers with an alternative approach for evaluating deep learning models and helps better visualise the ‘learning’ of these complex models. Vowel space also paves the way to develop deep learning approaches that can adjust the model’s training to learn the vowel space accurately. This study was conducted in the hope that future speech research will consider vowel space accuracy more closely when developing text to speech models. The work conducted in this study will give rise to better quality synthetic voices grounded on linguistics knowledge.

6. Acknowledgements

The authors thank the participants of the perceptions test and the University of Auckland Faculty of Engineering for funding.

³<https://github.com/babe269/performant>

7. References

- [1] C. Benoît, M. Grice, and V. Hazan, "The SUS test: A method for the assessment of text-to-speech synthesis intelligibility using Semantically Unpredictable Sentences," *Speech Communication*, vol. 18, no. 4, pp. 381–392, 1996.
- [2] P. C. Loizou, "Speech quality assessment," in *Multimedia analysis, processing and communications*. Springer, 2011, pp. 623–654.
- [3] A. Valizada, S. Jafarova, E. Sultanov, and S. Rustamov, "Development and evaluation of speech synthesis system based on deep learning models," *Symmetry*, vol. 13, no. 5, p. 819, 2021.
- [4] J. Kominek, T. Schultz, and A. W. Black, "Synthesizer voice quality of new languages calibrated with mean mel cepstral distortion," in *SLTU*, 2008, pp. 63–68.
- [5] S. Sandoval, V. Berisha, R. L. Utianski, J. M. Liss, and A. Spanias, "Automatic assessment of vowel space area," *The Journal of the Acoustical Society of America*, vol. 134, no. 5, pp. EL477–EL483, 2013.
- [6] A. Jongman, M. Fourakis, and J. A. Sereno, "The acoustic vowel space of Modern Greek and German," *Language and Speech*, vol. 32, pp. 221–248, 1989.
- [7] E. Jacewicz, R. A. Fox, and J. Salmons, "Vowel space areas across dialects and gender," in *Proc. Intl. Congress of Phonetic Sciences, Saarbrücken*, 2007, pp. 1465–1468.
- [8] C. G. Clopper and J. B. Pierrehumbert, "Effects of semantic predictability and regional dialect on vowel space reduction," *The Journal of the Acoustical Society of America*, vol. 124, no. 3, pp. 1682–1688, 2008.
- [9] C. Watson, B. Ross, E. Ballard, H. Charters, R. Arnold, and M. Meyerhoff, "Preliminary investigation into sound change in Auckland," in *Proc. Australasian Intl. Conf. on Speech Science and Technology*, 2018, pp. 17–20.
- [10] C. I. Watson, P. Keegan, M. Maclagan, R. Harlow, and J. King, "The motivation and development of MPAi, a Māori Pronunciation Aid," in *Proc. Interspeech*, 2017, pp. 2063–2067.
- [11] X. Chao, C. El-Khaissi, N. Kuo, P. K. John, and H. Suominen, "An automatic vowel space generator for language learner pronunciation acquisition and correction," in *Proc. Australasian Language Technology Association*, 2020, pp. 54–64.
- [12] A. Dowd, J. Smith, and J. Wolfe, "Learning to pronounce vowel sounds in a foreign language using acoustic measurements of the vocal tract as feedback in real time," *Language and Speech*, vol. 41, no. 1, pp. 1–20, 1998.
- [13] D. Brett, "Computer generated feedback on vowel production by learners of english as a second language," *ReCALL*, vol. 16, no. 1, pp. 103–113, 2004.
- [14] M. Carey, "Call visual feedback for pronunciation of vowels: Kay sona-match," *CALICO Journal*, pp. 571–601, 2004.
- [15] ANSI, "Acoustical Society of America standards," *Definition for formant*, 2004.
- [16] C. I. Watson, J. Harrington, and Z. Evans, "An acoustic comparison between New Zealand and Australian English vowels," *Australian Journal of Linguistics*, vol. 18, no. 2, pp. 185–207, 1998.
- [17] C. I. Watson and C. T. J. Hui, "Two short studies in vocal tract measurements," in *Proc. 13th Australasian International Conference on Speech Science and Technology*, pp. 9–12, 2010.
- [18] J. Hillenbrand, L. A. Getty, M. J. Clark, and K. Wheeler, "Acoustic characteristics of American English vowels," *The Journal of the Acoustical Society of America*, vol. 97, no. 5, pp. 3099–3111, 1995.
- [19] L. Albuquerque, C. Oliveira, A. J. Teixeira, P. Sa-Couto, and D. Figueiredo, "Age-related changes in European Portuguese vowel acoustics," in *Proc. Interspeech*, 2019, pp. 3965–3969.
- [20] R. Tatman and C. Kasten, "Effects of talker dialect, gender & race on accuracy of Bing speech and YouTube automatic captions," in *Proc. Interspeech*, 2017, pp. 934–938.
- [21] R. J. van Son, C. Middag, and K. Demuyne, "Vowel space as a tool to evaluate articulation problems," in *Proc. of the Intl. Speech Communication Association*, 2018, pp. 357–361.
- [22] G. S. Turner, K. Tjaden, and G. Weismer, "The influence of speaking rate on vowel space and speech intelligibility for individuals with amyotrophic lateral sclerosis," *Journal of Speech, Language, and Hearing Research*, vol. 38, no. 5, pp. 1001–1013, 1995.
- [23] J. C. Wells and J. C. Wells, *Accents of English: Volume 1*. Cambridge University Press, 1982, vol. 1.
- [24] J. C. Wells, *Accents of English: Volume 3: Beyond the British Isles*. Cambridge University Press, 1982, vol. 3.
- [25] K. Ito and L. Johnson, "The LJ speech dataset," <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [26] C. I. Watson and A. Marchi, "Resources created for building New Zealand English voices," in *Proc. Australasia. Intl. Conf. Speech Science and Technology*. Corpus available on request from c.watson@auckland.ac.nz, 2014, pp. 92–95.
- [27] C. I. Watson, "Mappings between vocal tract area functions, vocal tract resonances and speech formants for multiple speakers," in *Proc. of the Intl. Speech Communication Association*, 2014, pp. 1993–1997.
- [28] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan, R. A. Saurous, Y. Agiomvrgiannakis, and Y. Wu, "Natural TTS synthesis by conditioning Wavenet on MEL spectrogram predictions," in *Proc. Intl. Conf. on Acoustics, Speech Signal Processing*, 2018, pp. 4779–4783.
- [29] Y. Wang, R. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio, Q. Le, Y. Agiomvrgiannakis, R. Clark, and R. A. Saurous, "Tacotron: Towards end-to-end speech synthesis," in *Proc. Interspeech*, 2017, pp. 4006–4010.
- [30] A. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," in *Proc. SSW*, vol. 125, p. 2, 2016.
- [31] A. Gibiansky, S. Arik, G. Diamos, J. Miller, K. Peng, W. Ping, J. Raiman, and Y. Zhou, "Deep voice 2: Multi-speaker neural text-to-speech," *Advances in neural info. proc. sys*, vol. 30, 2017.
- [32] A. Debnath, S. S. Patil, G. Nadiger, and R. A. Ganesan, "Low-resource end-to-end Sanskrit TTS using Tacotron2, WaveGlow and Transfer Learning," in *Proc. India Council Intl. Conf. IEEE*, 2020, pp. 1–5.
- [33] G. Săracu and A. Stan, "An analysis of the data efficiency in Tacotron2 speech synthesis system," in *Proc. Intl. Conf. on Speech Technology and Human-Computer Dialogue (SpeD)*. IEEE, 2021, pp. 172–176.
- [34] "TensorFlowTTS," <https://github.com/TensorSpeech/TensorFlowTTS>, 2019.
- [35] Google, "Google colabouratory," 2017. [Online]. Available: <https://colab.research.google.com/>
- [36] F. Schiel, "Automatic Phonetic Transcription of Non-Prompted Speech," in *Proc. of the ICPHS*, August 1999, pp. 607–610.
- [37] D. Boersma, Paul Weenink, "Praat: doing phonetics by computer [computer program]," in *Version 6.0.46, 3 January 2019 from http://www.praat.org/*, 2019.
- [38] P. T. Inc. (2015) Collaborative data science. Montreal, QC. [Online]. Available: <https://plot.ly>
- [39] W. Labov, S. Ash, and C. Boberg, *The atlas of North American English: Phonetics, phonology and sound change*. Walter de Gruyter, 2008.