

Voice Conversion for Whispered Speech Synthesis

Marius Cotescu, Thomas Drugman, *Member, IEEE*, Goeric Huybrechts, Jaime Lorenzo-Trueba, *Member, IEEE*, and Alexis Moinet

Abstract—We present an approach to synthesize whisper by applying a handcrafted signal processing recipe and Voice Conversion (VC) techniques to convert normally phonated speech to whispered speech. We investigate using Gaussian Mixture Models (GMM) and Deep Neural Networks (DNN) to model the mapping between acoustic features of normal speech and those of whispered speech. We evaluate naturalness and speaker similarity of the converted whisper on an internal corpus and on the publicly available wTIMIT corpus. We show that applying VC techniques is significantly better than using rule-based signal processing methods and it achieves results that are indistinguishable from copy-synthesis of natural whisper recordings. We investigate the ability of the DNN model to generalize on unseen speakers, when trained with data from multiple speakers. We show that excluding the target speaker from the training set has little or no impact on the perceived naturalness and speaker similarity of the converted whisper. The proposed DNN method is used in the newly released Whisper Mode of Amazon Alexa.

Index Terms—whispered speech conversion, voice conversion (VC), whispered text to speech (TTS).

I. INTRODUCTION

WHISPER is a natural way of speaking, primarily characterized by the lack of phonation [1]. As Text-to-Speech (TTS) systems and virtual assistants become more natural and widespread, the scenarios when users need to interact with the devices in a quiet or private manner increase. In these cases, a user may wish to whisper to the device and would expect a response in a whispered voice, as is the case of the recently released Whisper Mode of Amazon Alexa [2]. While the generation of normally phonated speech from text (i.e. TTS), from the speech of other speakers (i.e. voice conversion), or even from whispered speech (e.g. pathological applications) is well researched, there has been little effort in developing technologies that can produce whispered speech.

In this paper, we aim at converting normal speech into whispered speech while maintaining high naturalness and speaker identity. To the best of our knowledge, this is the first study attempting to achieve this goal. For this purpose, two main approaches are investigated. First, we create a handcrafted digital signal processing (DSP) recipe that simulates acoustic modifications in a source-filter based vocoder [3], [4]. Second, we learn the mapping between normal and whispered acoustic features using voice conversion (VC) techniques.

Physiologically speaking, whispered speech differs from normal speech mainly by the lack of vocal folds vibration. Exhaled air passes directly through the restricted but open larynx which causes a turbulent aperiodic glottal airflow [1]. The acoustic manifestations of whispered speech are however

more diverse than a simple lack of pitch and harmonics. Extensive studies of the acoustic characteristics of whispered speech [5]–[9] found consistent changes in the formant structure, spectral tilt and consonant energy in whispered speech when compared to normal speech. Some of these findings were used to develop whispered speech detectors [10], [11]. We also draw on these findings to build a rule based DSP recipe that transforms normal speech into whispered speech.

VC models have traditionally relied on GMMs [12]–[14] to learn the joint distribution of acoustic features for source and target speakers from parallel corpora. More recently, feed forward, recurrent, and convolutional neural networks approaches [15]–[18] take advantage of larger parallel speech corpora to produce high-quality conversions. Other techniques achieve voice conversion by learning common latent [19]–[23] or explicit distributions [24] from non-parallel corpora. The task of converting whispered speech to normal speech has received significant research interest. There are many examples of applying VC techniques to the problem that vary from code excited linear prediction to recurrent neural networks [25], [26] and we draw inspiration from them.

The main contributions of this paper are: *i)* we explore the problem of normal to whispered speech conversion; *ii)* we propose a rule-based DSP system; *iii)* we use GMM and DNN models as VC techniques to learn the more complex mapping functions; *iv)* we evaluate our proposed solutions both on an internal proprietary corpus and on the publicly available wTIMIT¹ corpus [27]; *v)* we show that VC models can produce converted whisper of the same quality as vocoded natural recordings of whisper; *vi)* we show that through multi-speaker training, DNN models have the ability to generalize to unseen speakers.

II. DSP-BASED CONVERSION SYSTEM

We implemented a rule-based DSP transformation that applies the main acoustic modifications of whispered speech in the paradigm of the source-filter model. As vocoder, we opted for WORLD [4] which is widely used in speech synthesis for its high quality and computational efficiency. In addition to the obvious unvoiced setting, three main spectral changes were applied: 1) removing the glottal shaping, 2) shifting of the first formant, 3) increasing the formant bandwidth. *Step 1* removes the spectral coloring introduced by the voiced glottal air flow [28], [29] by subtracting the typical spectrum of a glottal pulse based on the LF model [30] from the spectral envelope. In *Step 2* we follow the findings reported in [6] and we use frequency warping of the spectral envelope to shift the first formant upwards by about 100 Hz while keeping the second and third for-

M. Cotescu, T. Drugman, G. Huybrechts, J. Lorenzo-Trueba and A. Moinet are with Amazon.com (e-mail: cotescu@amazon.com)
Manuscript received December 11, 2019.

¹The wTIMIT corpus is publicly available and can be downloaded from <http://www.isle.illinois.edu/sst/data/wTIMIT/>.

mants unchanged. In *Step 3* we apply a moving average filtering to the log spectrum with a 400 Hz-wide triangular window to mimic the increased formant bandwidth reported in [6], [9].

Although the DSP-based system may look simple, it can generalize on a diversity of male, female and child speakers. It was released in 2017 to support the Amazon Polly ‘whispered’ SSML tag [31] and is applied to 59 voices in 29 languages.

III. VOICE CONVERSION SYSTEM

We propose a conversion system based on the architecture presented in [32]. We extend it with the ability to use a feed-forward DNN conversion model and to train on multiple source and target speakers. A similar method was also reported in [33]. Our approach differs in that we apply it to acoustic data of multiple speakers to convert normal speech to whispered speech.

We start with one parallel corpus containing normal speech and whispered recordings of the same texts. Each pair of files is analyzed and acoustic features are extracted. The acoustic features frames are aligned using Dynamic Time Warping (DTW) in the iterative fashion described in [32]. The aligned features are used to train either a DNN or a GMM conversion model. At inference time, the source acoustic features are extracted from the input waveform and transformed by the conversion model. The converted features are then synthesized with the WORLD [4] vocoder.

A. Feature Extraction

Some VC methods choose to have separate feature sets and models for the source and vocal tract features. We choose to use a single set of features because *i*) whispered speech does not have pitch and *ii*) the voiced source characteristics inform the spectral envelope transformations. We use 80 mel-frequency cepstral coefficients (MFCC) [34]. The cepstral domain provides a compact and untangled representation of both the vocal tract and the excitation, where the lower order coefficients are related to the vocal tract shape and the higher order coefficients describe excitation source traits (e.g. voicing strength, pitch) [35]. The cepstral coefficients are computed on 50 ms Hanning analysis windows, with a frame shift of 5 ms. All recordings were down-sampled to 24 kHz.

B. GMM Whisper Conversion

We used *sprocket* [32], an open-source implementation of GMM-based voice conversion [13]. It was used as a baseline in the 2018 Voice Conversion Challenge [36] and it thus gives us a solid anchor when comparing to other VC methods. We tuned the hyper parameters of the GMM conversion model on the internal validation set (see Section IV-A). We observed the best results using a 64-mixture GMM with the convergence threshold set at 0.001. At inference time, we used Maximum Likelihood Parameter Generation (MLPG) considering the Global Variance (GV) [37] to generate the converted target cepstral coefficients.

C. DNN Whisper Conversion

In this paper, we propose using a DNN model to learn a general mapping from acoustic features of normal speech to those of whispered speech. Neural networks can be trained on batches [38] and benefit from several regularization

TABLE I
TECHNOLOGY AND DATA USED FOR THE EVALUATED SYSTEMS

Name	Tech.	Data
Rec	-	Natural whisper recordings
Oracle	Vocoder	Vocoded whisper recordings
sprocket	GMM	Only the evaluated speaker
SD	DNN	Only the evaluated speaker
All	DNN	The entire internal or wTIMIT corpus
Excl	DNN	All speakers in the internal or wTIMIT corpus, excluding the evaluated speaker
DSP	DSP	-

techniques [39]–[41] that allow them to generalize better on large diverse training sets.

We use a neural network with 4 hidden layers, of size (128, 64, 64, 128) with a rectified linear activation function. We use z-score normalization for both the input and output feature vectors. At inference time, the estimated target vector is scaled back, before being sent to the processing chain. We train the network using the Adam [42] optimizer to minimize the L2 loss. We use a batch size of 2048 frames, a learning rate of 0.002, and L2 weights regularization [39], with $\lambda = 10^{-5}$. The hyper parameters were tuned on the internal validation set, described in Section IV-A.

IV. EXPERIMENTAL PROTOCOL

A. Datasets

We used two corpora to train and evaluate our systems. The first one is an internal corpus of parallel whispered and normal speech recordings. It contains recordings of phonetically rich texts from 5 English female speakers, each from a different locale: Australia (AU), Canada (CA), Great Britain (GB), India (IN) and United States (US). All speakers are professional voice talents and the recordings were done with high quality equipment. The corpus contains 10 hours of both normal and whispered speech (AU 2h, CA 3h, GB 1.5h, IN 3h, US 0.5h).

The second dataset is a subset of the wTIMIT [27] corpus. It contains parallel recordings of normal and whispered speech from 50 male and female English speakers from Singapore (SG) and the United States (US). The recordings average at 50 minutes per speaker. The SNR and quality of the recordings varies greatly between speakers. To address this problem, we manually selected 25 speakers for which the recordings are decently clean. The clean subset contains 20 US speakers and 5 SG speakers (13 male and 12 female).

For both datasets, we randomly split the data of each speaker into a training, a validation and a test set. We used a 80/10/10% split, respectively.

B. Systems

Table I shows the technologies and the data used to train the different conversion models. The **Oracle** system consists of copy-synthesis samples reconstructed with the WORLD [4] from cepstral coefficients computed on original whisper recordings. It constitutes a strong baseline to which we report all our proposed methods. We trained speaker-dependent DNN (**SD**) and GMM (**sprocket**) models for each speaker in the two datasets. The **All** models were trained on all the speakers from a given dataset (i.e. either wTIMIT or the internal one), while the **Excl** models were trained on all the speakers in the dataset, except the target speaker. The **DSP** samples are synthesized using the technique described in Section II.

C. Listening Tests

We used naturalness listening tests to measure the general quality of the conversion and speaker similarity tests to confirm that the speaker identity is maintained. We used self-reported effort to measure the intelligibility of the proposed systems. We used modified Multiple Stimuli with Hidden Reference and Anchor (MUSHRA) [43] for the naturalness and similarity tests and a Mean Opinion Score (MOS) for the intelligibility test. We modified the MUSHRA tests by not enforcing that at least one system gets the maximum rating.

For the naturalness tests, we asked the listeners to evaluate the stimuli in terms of their naturalness on a scale from 0 – “Completely unnatural” to 100 – “Completely natural”. In the speaker similarity tests, in addition to the evaluated stimuli, we also presented the listeners with a reference recording of the target speaker whispering a different phrase. We asked the listeners to score the speaker similarity of the stimuli to the provided reference on a scale from 0 – “A completely different speaker” to 100 – “The same speaker”. In the intelligibility test we asked the listeners to evaluate the effort they needed to understand the message. They reported on a scale from 1 “Very difficult: I cannot understand what was said” to 7 “Very easy: I understood everything”. We suggested 4 as the middle point “Fair: I understood everything, but I had to focus”.

For each speaker we randomly chose 50 phrases from the test set, which were used in all the listening tests. We recruited listeners on the Amazon Mechanical Turk platform and each listener evaluated exactly 10 phrases. We had no guarantee of their English proficiency, or the equipment they used. Following the guidelines in [44] we use Intra-Class Correlation (ICC(2,k)) [45] to measure the reliability of our results. The test reliability is said to be good if the ICC value is above 0.75 and excellent if greater than 0.90 [44]. We present the value in all boxplots. We use paired Student t-tests with Holm-Bonferroni correction to validate the statistical significance of the differences between systems, considering it validated when $p\text{-value} < 0.05$.

V. RESULTS

A. General Quality

We assessed the general quality of the proposed techniques by comparing them to the copy-synthesis system **Oracle**. We conducted naturalness MUSHRA and intelligibility MOS listening tests for the US internal speaker. We asked 100 listeners to evaluate samples from five systems: **Rec**, **Oracle**, **SD**, **sprocket**, and **DSP**. In the intelligibility test we included samples of normal speech synthesized using the Amazon Polly service (**TTS**) as a high anchor. Fig. 1 shows the boxplots for the two listening tests.

Both VC methods perform much better than the handcrafted DSP method and achieve the maximum expected naturalness and intelligibility. They performed better or equal to the Oracle system, reaching the limits of the technology. In single speaker scenarios one can use any of the two methods with equal results.

Listeners find it harder to understand whispered speech than they do normal speech [46]. This is demonstrated by the big difference between the **TTS** system and the whispered




Fig. 1. Naturalness (a) and intelligibility (b) scores for the US voice. Horizontal bars link systems for which $p\text{-value} > 0.05$.

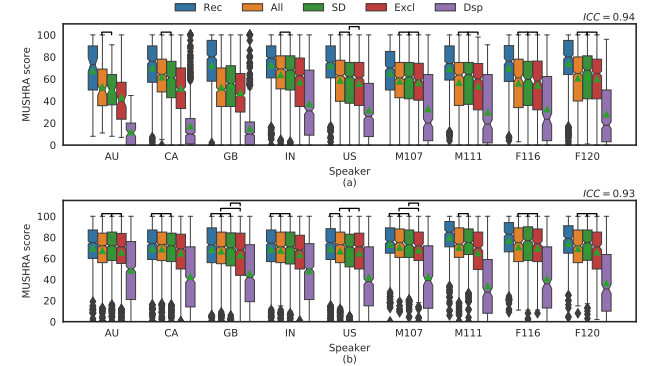


Fig. 2. Intra-dataset naturalness (a) and speaker similarity (b) scores. Horizontal bars link systems for which $p\text{-value} > 0.05$.

recordings. However, both VC methods are consistently scoring above 4.

B. Multi-Speaker Conversion

In this section we evaluate the ability of the DNN model to generalize to unseen speakers when trained on data from multiple speakers. Although both the DNN and the GMM systems performed similarly in the speaker independent task, we could not train GMM models on multi-speaker datasets. Multiple speakers datasets contained unpredictable modalities that the GMM model could not capture well. We tried to apply **SD** models to unseen speakers and had unstable results so we will not present evaluations of this scenario either.

We conducted naturalness and speaker similarity MUSHRA tests on both the internal corpus and the wTIMIT corpus. We observed the same quality patterns for all wTIMIT speakers in informal listening tests, independent of the recording conditions or accent. However, for the formal listening tests we chose two male (M107, M111) and two female (F116, F120) speakers with clean recording environments for both normal and whispered sets.

Intra-Dataset Generalization: We evaluated samples from five systems: **Rec**, **All**, **SD**, **Excl** and **DSP**. The **All** and **Excl** models were trained only on speakers from the same dataset as the evaluated speaker. We used 100 listeners for each test on internal speakers and 40 listeners for the wTIMIT speakers. Fig. 2 shows the boxplots of the MUSHRA scores in the naturalness and speaker similarity tests for the internal and wTIMIT speakers.

The ranking of the systems’ average naturalness and speaker similarity scores is consistent across most of the speakers in both datasets. The reference system **Rec** is getting the highest scores, followed by a close grouping of the multi-speaker (**All**) and the speaker dependent (**SD**) systems. The differences in naturalness and speaker similarity between the two models are not statistically significant, except in the case of the internal GB speaker. The multi-speaker (**Excl**)




Fig. 3. Cross-dataset naturalness (a) and speaker similarity (b) scores. Horizontal bars link systems for which p -value > 0.05 .

systems excluding the target speaker follow closely. We observe that removing the target speaker from the training set has little effect in the internal corpus and almost no effect in the wTIMIT corpus. The naturalness and speaker similarity scores of the **Excl** system are slightly lower than those of **All**. There is no statistically significant difference in naturalness for five speakers and in speaker similarity for six speakers. The **DSP** system has the lowest mean score in all tests.

Training a model with data from more speakers does not lead to a significant increase in quality on seen speakers when compared to speaker dependent models. This is consistent with the results from Section V-A. However, it has a positive impact on the model's ability to generalize to unseen speakers. The **Excl** models trained on internal speakers can produce whisper for unseen speakers with only a small drop in quality, independent of their accent, while the ones trained on the wTIMIT corpus are performing as well as the **All** and the **SD** models. We believe that the more varied recording conditions and speaker characteristics present in the wTIMIT corpus helps the models learn a more general transformation.

Cross-Dataset Generalization: We assessed the ability of the systems to generalize across the two datasets. We asked 30 listeners to evaluate the naturalness and speaker similarity of five systems: **Rec**, **All^{Int}**, **All^{wTIMIT}**, **Excl** and **DSP** systems. The **Excl** system was always trained on the dataset of the target speaker. Because there are no male speakers in the internal corpus, we only evaluated the two wTIMIT female speakers (F116, F120). We chose the IN and US internal speakers for formal evaluation as they performed the best in the intra-dataset evaluation. Informal listening tests confirmed the same quality pattern across all internal speakers. Fig. 3 shows the boxplots of the scores.

We find that the models need diverse data in order to generalize to more diverse conditions. Although the **All^{Int}** model is performing much better than the handcrafted DSP system on the wTIMIT speakers, it is still significantly worse than the **Excl** system trained on wTIMIT data. This is caused partly by the fact that the internal corpus only contains 5 speakers and partly by the fact that the wTIMIT recordings are noisier. The noise in the wTIMIT corpus also prevented the **All^{wTIMIT}** model from performing better on the internal corpus. However, the presence of more varied data helped it generalize even to unseen recording conditions and it performed on par with the **Excl** models on the internal speakers.

Cross-Gender Generalization: We assessed the ability of the DNN models to generalize across genders in the wTIMIT corpus. We asked 30 listeners to evaluate the naturalness and speaker similarity of five systems: **Rec**, **All**, **Female**, **Male** and **DSP** systems for the four wTIMIT evaluation speakers. The **All** model was trained on the entire wTIMIT corpus,




Fig. 4. Cross-gender naturalness (a) and speaker similarity (b) scores. Horizontal bars link systems for which p -value > 0.05 .

while the **Male** and **Female** models were trained on all the male and female wTIMIT speakers, respectively. Fig. 4 shows the boxplots of the MUSHRA scores.

The gender balance is less important when the target speaker is present in the training set: the matched gender models perform as well as the **All** model. However, the model seemed to need gender balanced data in order to generalize well to unseen speakers from both genders. Using models trained on either male or female speakers on unseen speakers of the other gender produced whisper in every case. However, the gender identity of the converted samples was affected and the quality was varying, as reflected by the lower means and higher variances of the naturalness and speaker similarity scores. This drop cannot be explained by the data reduction, as the cross-dataset generalization experiments show less degradation when applying the model trained on the internal corpus to the wTIMIT female speakers.

The poor performance in the cross-gender scenario is caused by the fact that the acoustic characteristics of the voiced excitation in the cepstral space vary significantly between the genders. They vary not only in pattern, but also in where they occur in the cepstrum (i.e. the inverse relationship between pitch and quefrequency). We highly recommend using gender balanced training datasets in applications where cross gender generalization is important.

VI. CONCLUSIONS

This paper was, to the best of our knowledge, the first study attempting to convert normal speech into whispered speech. We investigated using signal processing and two VC techniques (based on GMM and DNN models) to achieve this goal. We evaluated the three methods on both an internal corpus and on the wTIMIT public corpus.

We found out that the DSP recipe had a robust performance over all speakers from both corpora, however it is far from the theoretical limit imposed by the vocoder. All VC approaches however outperformed the DSP system in naturalness, intelligibility and speaker similarity evaluations. They reached the technical limit imposed by the vocoder and feature extraction chain. We showed that the DNN models can learn a speaker-independent mapping when trained on multiple speakers and that they are able to generalize and produce whispered speech for unseen speakers. We found that the DNN models can be robust to recording conditions if trained with varied enough data. The DNN model cannot generalize across genders and a gender balanced corpus is recommended for cross gender applications. The proposed DNN method was integrated into Amazon Alexa, and is used to generate the output for the newly released Whisper Mode [2].

REFERENCES

- [1] P. Monoson and W. R. Zemlin, "Quantitative study of whisper," *Folia Phoniatrica et Logopaedica*, vol. 36, no. 2, pp. 53–65, 1984.
- [2] "Alexa is now even smarter," <https://press.aboutamazon.com/news-releases/news-release-details/alexa-now-even-smarter-new-features-help-make-everyday-life-more>, 2018, accessed: 2019-03-13.
- [3] G. Fant, *Acoustic theory of speech production*. Walter de Gruyter, 1970, no. 2.
- [4] M. Morise, F. Yokomori, and K. Ozawa, "World: a vocoder-based high-quality speech synthesis system for real-time applications," *IEICE TRANSACTIONS on Information and Systems*, vol. 99, no. 7, pp. 1877–1884, 2016.
- [5] K. J. Kallail and F. W. Emanuel, "Formant-frequency differences between isolated whispered and phonated vowel samples produced by adult female subjects," *Journal of Speech, Language, and Hearing Research*, vol. 27, no. 2, pp. 245–251, 1984.
- [6] H. R. Sharifzadeh, I. V. McLoughlin, and M. J. Russell, "A comprehensive vowel space for whispered speech," *Journal of voice*, vol. 26, no. 2, pp. e49–e56, 2012.
- [7] S. T. Jovičić and Z. Šarić, "Acoustic analysis of consonants in whispered speech," *Journal of voice*, vol. 22, no. 3, pp. 263–274, 2008.
- [8] T. Itoh, K. Takeda, and F. Itakura, "Acoustic analysis and recognition of whispered speech," in *IEEE Workshop on Automatic Speech Recognition and Understanding, 2001. ASRU'01*. IEEE, 2001, pp. 429–432.
- [9] V. C. Tarter, "What's in a whisper?" *The Journal of the Acoustical Society of America*, vol. 86, no. 5, pp. 1678–1683, 1989.
- [10] Z. Raeesy, K. Gillespie, C. Ma, T. Drugman, J. Gu, R. Maas, A. Rastrow, and B. Hoffmeister, "Lstm-based whisper detection," in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 139–144.
- [11] C. Zhang and J. H. Hansen, "Advancements in whisper-island detection within normally phonated audio streams," in *Tenth Annual Conference of the International Speech Communication Association*, 2009.
- [12] Y. Stylianou, O. Cappé, and E. Moulines, "Continuous probabilistic transform for voice conversion," *IEEE Transactions on speech and audio processing*, vol. 6, no. 2, pp. 131–142, 1998.
- [13] T. Toda, A. W. Black, and K. Tokuda, "Voice conversion based on maximum-likelihood estimation of spectral parameter trajectory," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 8, pp. 2222–2235, 2007.
- [14] E. Helander, T. Virtanen, J. Nurminen, and M. Gabbouj, "Voice conversion using partial least squares regression," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 5, pp. 912–921, 2010.
- [15] S. H. Mohammadi and A. Kain, "Voice conversion using deep neural networks with speaker-independent pre-training," in *2014 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2014, pp. 19–23.
- [16] Y. Saito, S. Takamichi, and H. Saruwatari, "Voice conversion using input-to-output highway networks," *IEICE Transactions on Information and Systems*, vol. 100, no. 8, pp. 1925–1928, 2017.
- [17] L. Sun, S. Kang, K. Li, and H. Meng, "Voice conversion using deep bidirectional long short-term memory based recurrent neural networks," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 4869–4873.
- [18] K. Tanaka, H. Kameoka, T. Kaneko, and N. Hojo, "AttS2S-VC: Sequence-to-sequence voice conversion with attention and context preservation mechanisms," in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 6805–6809.
- [19] C.-C. Hsu, H.-T. Hwang, Y.-C. Wu, Y. Tsao, and H.-M. Wang, "Voice conversion from unaligned corpora using variational autoencoding wasserstein generative adversarial networks," in *Proc. Interspeech 2017*, 2017, pp. 3364–3368.
- [20] H. Kameoka, T. Kaneko, K. Tanaka, and N. Hojo, "ACVAE-VC: Non-parallel many-to-many voice conversion with auxiliary classifier variational autoencoder," *arXiv preprint arXiv:1808.05092*, 2018.
- [21] W. Huang, Y. Wu, H. Hwang, P. L. Tobing, T. Hayashi, K. Kobayashi, T. Toda, Y. Tsao, and H. Wang, "Refined wavenet vocoder for variational autoencoder based voice conversion," in *2019 27th European Signal Processing Conference (EUSIPCO)*, Sep. 2019, pp. 1–5.
- [22] F. Fang, J. Yamagishi, I. Echizen, and J. Lorenzo-Trueba, "High-quality nonparallel voice conversion based on cycle-consistent adversarial network," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5279–5283.
- [23] H. Kameoka, T. Kaneko, K. Tanaka, and N. Hojo, "StarGAN-VC: Non-parallel many-to-many voice conversion with star generative adversarial networks," *arXiv preprint arXiv:1806.02169*, 2018.
- [24] X. Tian, J. Wang, H. Xu, E.-S. Chng, and H. Li, "Average modeling approach to voice conversion with non-parallel data," in *Proc. Odyssey 2018 The Speaker and Language Recognition Workshop*, 2018, pp. 227–232.
- [25] H. R. Sharifzadeh, I. V. McLoughlin, and F. Ahmadi, "Reconstruction of normal sounding speech for laryngectomy patients through a modified CELP codec," *IEEE Transactions on Biomedical Engineering*, vol. 57, no. 10, pp. 2448–2458, 2010.
- [26] G. N. Meenakshi and P. K. Ghosh, "Whispered speech to neutral speech conversion using bidirectional lstms," in *Proc. Interspeech 2018*, 2018, pp. 491–495.
- [27] B. P. Lim, "Computational differences between whispered and non-whispered speech," Ph.D. dissertation, University of Illinois at Urbana-Champaign, 2011.
- [28] B. Doval, C. d'Alessandro, and N. Henrich, "The spectrum of glottal flow models," *Acta acustica united with acustica*, vol. 92, no. 6, pp. 1026–1046, 2006.
- [29] T. Drugman, P. Alku, A. Alwan, and B. Yegnanarayana, "Glottal source processing: From analysis to applications," *Computer Speech & Language*, vol. 28, no. 5, pp. 1117–1138, 2014.
- [30] G. Fant, J. Liljencrants, and Q.-g. Lin, "A four-parameter model of glottal flow," *STL-QPSR*, vol. 4, no. 1985, pp. 1–13, 1985.
- [31] "Amazon 'whispered' ssm tag," <https://aws.amazon.com/about-aws/whats-new/2017/04/amazon-polly-introduces-a-new-whispered-voice-tag-and-speech-marks-for-synchroni>, 2017, accessed: 2019-03-13.
- [32] K. Kobayashi and T. Toda, "sprocket: Open-source voice conversion software," in *Proc. Odyssey 2018 The Speaker and Language Recognition Workshop*, 2018, pp. 203–210.
- [33] G. NishaMeenakshi and P. K. Ghosh, "Reconstruction of articulatory movements during neutral speech from those during whispered speech," *The Journal of the Acoustical Society of America*, vol. 143 6, p. 3352, 2018.
- [34] S. Imai, "Cepstral analysis synthesis on the mel frequency scale," in *ICASSP'83. IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 8. IEEE, 1983, pp. 93–96.
- [35] A. Oppenheim and R. Schaffer, "Homomorphic analysis of speech," *IEEE Transactions on Audio and Electroacoustics*, vol. 16, no. 2, pp. 221–226, 1968.
- [36] J. Lorenzo-Trueba, J. Yamagishi, T. Toda, D. Saito, F. Villavicencio, T. Kinnunen, and Z. Ling, "The Voice Conversion Challenge 2018: Promoting development of parallel and nonparallel methods," in *Proc. Odyssey 2018 The Speaker and Language Recognition Workshop*, 2018, pp. 195–202.
- [37] T. Toda and K. Tokuda, "A speech parameter generation algorithm considering global variance for hmm-based speech synthesis," *IEICE TRANSACTIONS on Information and Systems*, vol. 90, no. 5, pp. 816–824, 2007.
- [38] L. Bottou, "Stochastic gradient learning in neural networks," *Proceedings of Neuro-Nimes*, vol. 91, no. 8, p. 12, 1991.
- [39] A. Krogh and J. A. Hertz, "A simple weight decay can improve generalization," in *Advances in neural information processing systems*, 1992, pp. 950–957.
- [40] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [41] S. Han, J. Pool, J. Tran, and W. Dally, "Learning both weights and connections for efficient neural network," in *Advances in neural information processing systems*, 2015, pp. 1135–1143.
- [42] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [43] B. ITU-R Recommendation, "1534-1: Method for the subjective assessment of intermediate quality levels of coding systems (MUSHRA)," *International Telecommunication Union*, 2003.
- [44] T. K. Koo and M. Y. Li, "A guideline of selecting and reporting intraclass correlation coefficients for reliability research," *Journal of chiropractic medicine*, vol. 15, no. 2, pp. 155–163, 2016.
- [45] A. Donner and J. J. Koval, "The estimation of intraclass correlation in the analysis of family data," *Biometrics*, pp. 19–25, 1980.
- [46] R. L. Freyman, A. M. Griffin, and A. J. Oxenham, "Intelligibility of whispered speech in stationary and modulated noise maskers," *The Journal of the Acoustical Society of America*, vol. 132, no. 4, pp. 2514–2523, 2012.